Navigational Ability in Hyperbolic Space: A Study in VR

R.P.G. van Buren

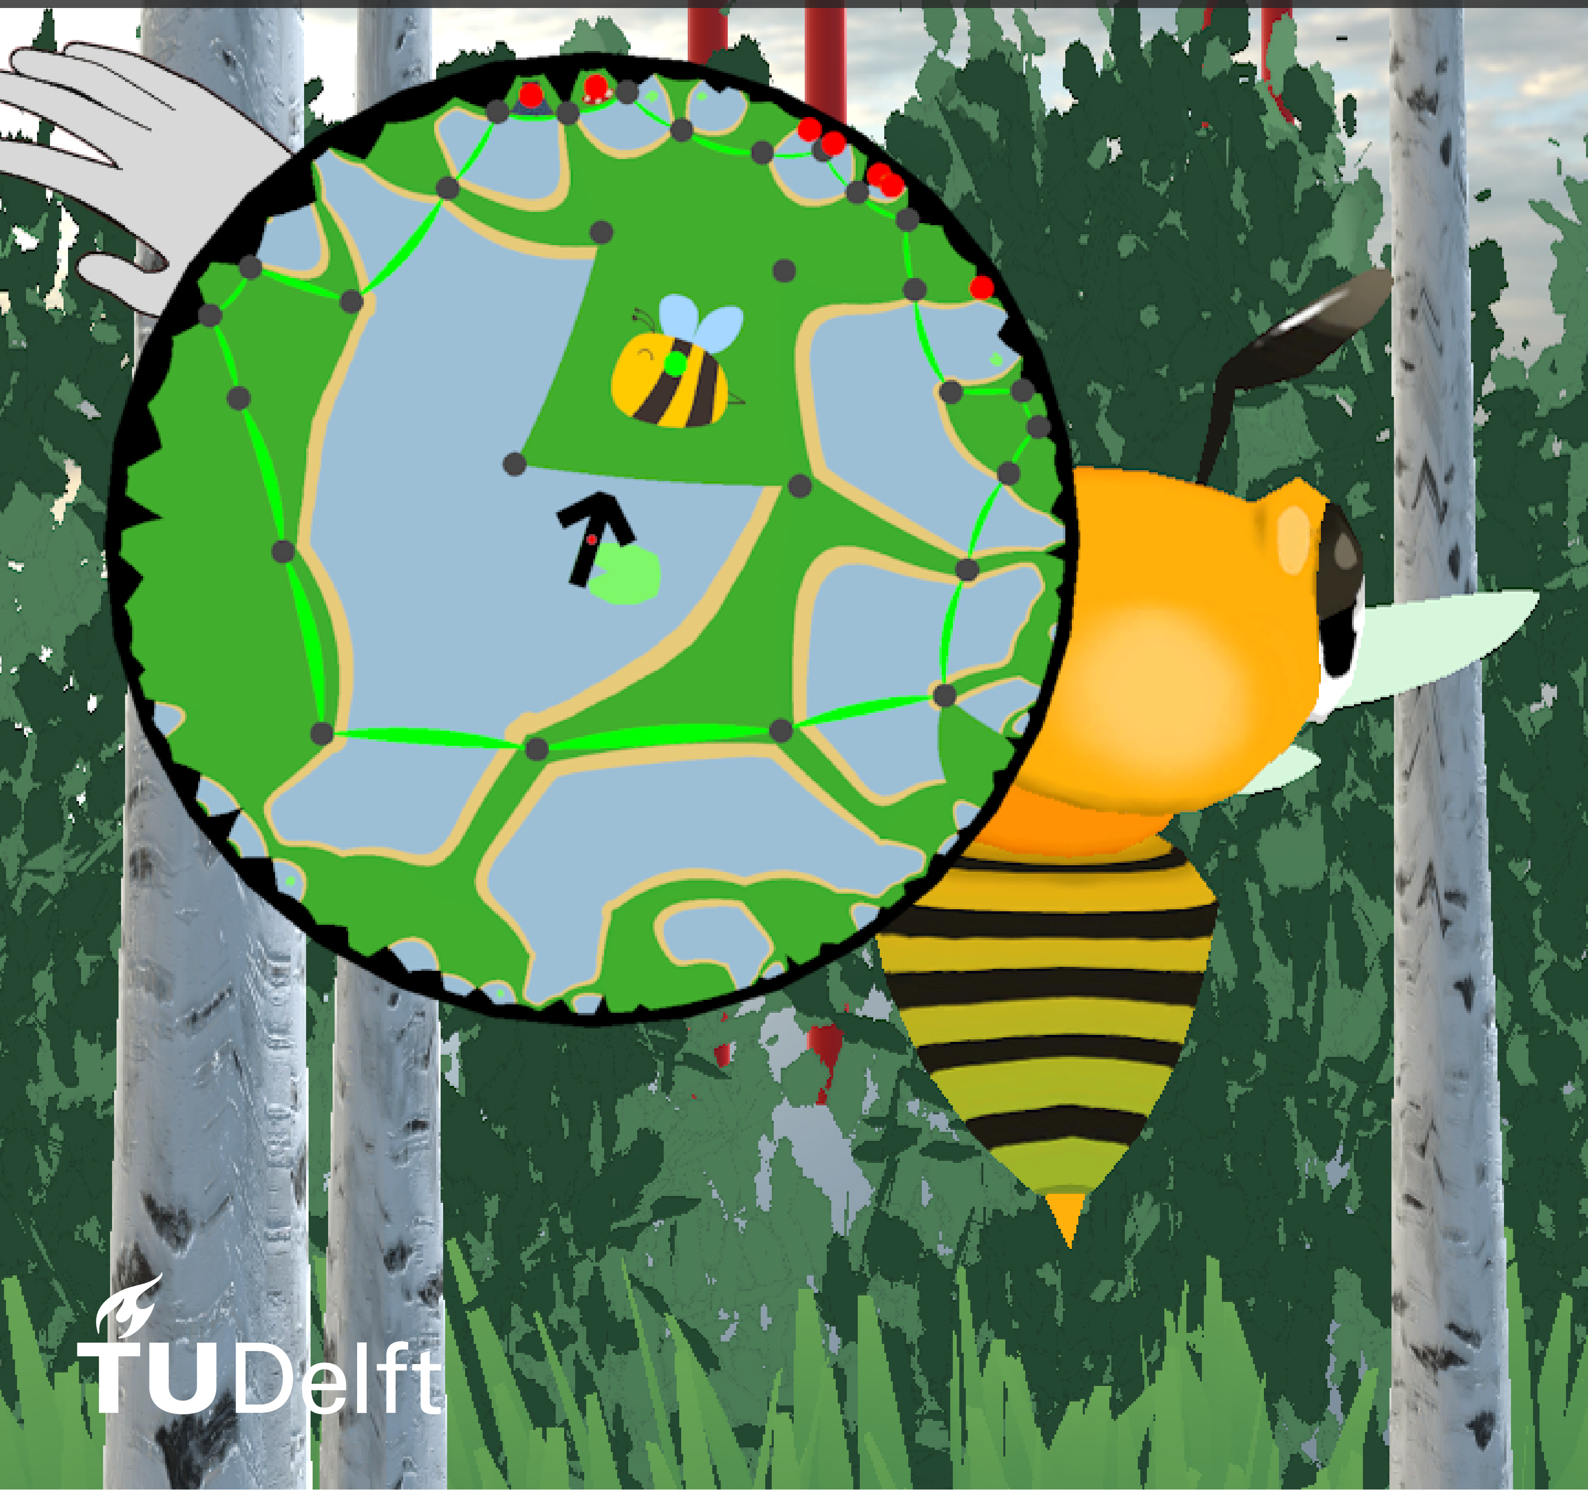

Abstract

With Virtual Reality, we can create and explore an infinite number of environments. These environments can have multiple applications, such as in education, training, or entertainment. However, we need a way to move through these environments. The most natural way is to walk, but we are limited by our physical space. A solution would be hyperbolic space. In this thesis, we explore navigational performance in hyperbolic space and its relation to the properties of hyperbolic geometry. By using the existing application Holonomy VR, we can use hyperbolic space and virtual reality to create an infinite world in a confined physical space. We let participants navigate in this virtual experience to study if they can navigate and learn about its properties. To achieve this, the performance of Holonomy VR is improved and new features are added that allow for a textured top view of the hyperbolic plane in the form of a 2D experience. Also, a new mode is added that requires users to find multiple landmarks in the environment. The beacon cue is developed to let users see targets outside of their current reach. A user study is conducted in collaboration with Leiden University to test the navigational performance and the understanding of the hyperbolic properties. The results show that some people can adapt to navigating in a confined hyperbolic space and that the 2D experience results in the fastest training time. However, it is found that training by exploring the environment and learning about the landmarks before the evaluation does not improve the navigational performance. This contradicts the findings about navigation in a Euclidean scene. It is also found that some hyperbolic properties can be learned and that the 2D experience made it easier to learn about the diverging paths property. Participants who grasped the concept found the experience fun and interesting, and some even indicated that they wanted to learn more about hyperbolic space. Overall, this study shows that people can navigate a confined hyperbolic environment and that they can learn about its properties.

Contents

1
Introduction

Have you ever gotten lost in a new place and did not know how to get back, or how to reach your goal? Well, you are probably not the only one. Daily we use our navigational skills to get from point A to point B. Doing the same trip often enough will make you familiar with the route you take and allow you to travel with ease. However, when you are in a new place you need to rely on your navigational skills and the cues that are available to you. These cues can be anything from a list of directions to your precise location displayed on your phone. How we navigate and how we improve our navigational skills have been researched thoroughly in the past by using real-world examples but also using virtual experiences and virtual reality.

Using virtual reality we can create fascinating environments and explore environments that are close to real life, or very far from it. However, when using virtual reality, people can experience motion sickness, anxiety, eye strain, headaches, and other symptoms which are generally known as cybersickness [48, 7, 58, 42]. This can have multiple causes ranging from the quality of the headset to the experience provided, but it is also a personal matter. One of the main causes of motion sickness is the illusion of motion [58, 36]. This is often necessary when we want to move around in the virtual space because we are still limited by the physical size of the room. To overcome this issue we can look at hyperbolic geometry. This allows us to create infinitely more space in the confined physical space. However, hyperbolic space can be confusing and possibly hard to navigate. But can people get familiar with this non-Euclidean space and learn to navigate in it? The answer is yes, and they can even learn about some of its properties. Since hyperbolic geometry is not often taught, this experience might excite people to learn more about it and explore its use cases.

In this thesis, we will explore the navigational performance of people in hyperbolic space. People are very flexible when navigating. However, they need to be familiar with the environment or use cues to find their way to the target. They can be familiarised by training with the environment which can also be done with specific cues. Using cues to find their target can teach them about the environment, but the cue should still promote active exploration of the environment. But, can we also train people to navigate efficiently in a hyperbolic scene? How does research about navigational performance in the Euclidean scene carry over to a hyperbolic scene? And can people learn some of the properties of hyperbolic space by exploring it? Before we can hypothesise these questions we need to understand the current state of research in the field of navigation and hyperbolic geometry. This is discussed in chapter 3. Then we state and hypothesise the research questions in chapter 4.

Answering the research questions allows us to get an insight into navigation in hyperbolic space and find a novel way to teach people about hyperbolic space. These results can be used to improve the design of virtual environments to use non-Euclidean geometries to fit more content into a limited physical space and to improve the locomotion experience in virtual reality (VR) applications and possibly reduce the risk of motion sickness. By letting the user navigate the virtual world by walking we create a natural way to explore a hyperbolic scene. By using the holonomy property of hyperbolic space we can overcome the boundaries of our physical space and create an infinite hyperbolic scene in this physically limited space. As one might expect this confinement makes it harder to explore the created scene since the user is required to exploit the properties of hyperbolic space to be able to reach their target. But most users should be able to adapt to these limitations. These concepts, and how this is already achieved in Holonomy VR [50], are discussed in chapter 2.

Holonomy VR provides a virtual environment (VE) that lets users explore the hyperbolic plane in VR. To improve the user experience the application is updated in multiple ways. First by enhancing the performance to allow it to run better on standalone VR headsets. Then the map is improved to show the user more information about the environment. Finally, a new experience is added to the application that requires the user to navigate to different landmarks. This mode will help us to understand how people navigate in hyperbolic space and to analyse their performance. These improvements and additions are discussed in chapter 5.

To validate the hypotheses we will conduct a user study, which is described in chapter 6. The user study is held together with Leiden University with the faculty of ”Social and Behavioural Sciences” who will focus on the psychological part of the navigational ability. The study is split into three parts, in the first part we collect data about the participant. Then we let the participant practice to get familiar with Holonomy VR and train them to familiarise them with the landmarks found in the VE. After the training we evaluate their performance by letting them find the landmarks in a specific order. In the last part of the survey we test their knowledge on the location of the landmarks and the properties of hyperbolic space by letting them answer questions about the locations of the landmarks and the properties of the used hyperbolic space.

The results of the study are discussed in chapter 7 and analysed in chapter 8. Here it is found that not everyone was able to navigate successfully, but people can definitely adapt to exploring a confined hyperbolic space. However, training them by letting them first find all the available landmarks did not improve navigational performance. But, people who got to experience Holonomy VR in 2D by using only a simplified top view did learn quicker. People were also able to relate their experience to some of the hyperbolic properties. The people that grasped the concept found the experience interesting and fun. These conclusions are explained in more detail in chapter 9.

In summary, the contributions of this thesis are:

- Technical improvements to Holonomy VR that allow users to explore the hyperbolic plane in 2D.
- Addition of a beacon cue that allows users to see targets outside of their current reach.
- Developed a new mode in Holonomy VR that requires users to navigate to multiple landmarks.
- Compared different training methods to see if they improve navigational performance.
- Conducted a user study to evaluate the navigational performance of users in a confined hyperbolic space and their understanding of the hyperbolic properties.

2 Background

In this chapter we lay the basis to understand the research. We start by explaining non-Euclidean geometry and how it differs from Euclidean geometry. We then introduce Holonomy VR, the application that forms the foundation of this thesis and which we will expand upon.

2.1. Non-Euclidean Geometry

Non-Euclidean geometry is a type of geometry that fails to satisfy the 5th postulate of Euclid, the parallel postulate. This postulate states that if two straight lines intersect another line with 2 interior angles on the same side that are less then 90 degrees, then these two lines will intersect at some point. Breaking this postulate allows for the creation of geometries that are different from the familiar Euclidean geometry. There are two main types of non-Euclidean geometry: spherical geometry and hyperbolic geometry. In this thesis, we focus on hyperbolic geometry. Nevertheless, it is good to have an introduction to spherical geometry as well since it provides us with an easier to understand example of non-Euclidean geometry. Figure 2.1 shows an overview of the properties of these types of non-Euclidean geometry.

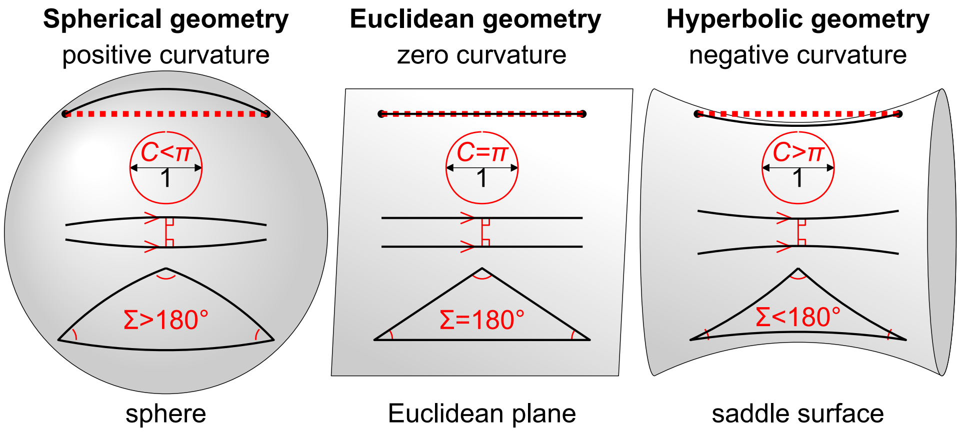

Figure 2.1: Comparison between spherical, Euclidean and hyperbolic geometry [9]

Spherical geometry is a type of non-Euclidean geometry that is based on the surface of a sphere. In this geometry, the sum of the angles of a triangle is greater than 180 degrees, and parallel lines do not exist because every line intersects another line. This means that the shortest distance between two points is not a straight line, but a curve known as a great circle. Spherical geometry is used to model the surface of the Earth.

Since the geometry is a sphere, every line on this sphere has become a curve. This results in interesting behaviour when using parallel transport. Parallel transport involves moving a vector along a curve on the sphere while keeping the vector parallel to itself according to the sphere's curvature. In non-Euclidean geometry this process will change the direction of the vector due to the curvature of the space. To understand parallel transport, consider the example below as shown in Figure 2.2.

This can be applied to our own world. Imagine that you are standing on the equator and looking north. You start walking north and you keep walking in a straight line until you reach the North Pole. Now, without turning, you will move to the right until you are back at the equator. Finally, walk backwards along the equator until you are back at the point where you started. You will notice that you are now facing east while you started facing north. This change in direction is the result of parallel transport and is called holonomy. As you moved along the curved surface of the Earth, the direction of your vector (initially pointing north) changed due to the curvature of the space. This effect is also present in hyperbolic geometry and is used to rotate the environment in Holonomy VR.

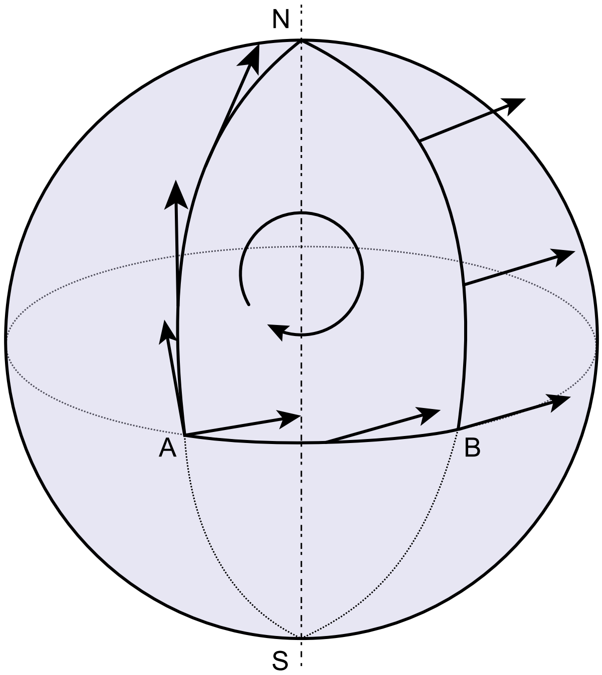

Figure 2.2: Moving from A to N to B and back to A causes the direction of the initial vector to rotate 90 degrees clockwise [54]

Hyperbolic geometry is another type of non-Euclidean geometry. In hyperbolic geometry, the sum of the angles of a triangle is less than 180 degrees, and parallel lines diverge from each other. This means that the shortest distance between two points is not a straight line, but a curve known as a hyperbolic line. There are multiple ways to model the hyperbolic plane. In this thesis, we use two of these models, the Poincaré disk model and the hyperboloid model (or Minkowski model). In the Poincaré disk model, the hyperbolic plane is mapped onto the unit disk where the edge of the disk represents infinity. In the hyperboloid model, the hyperbolic plane is represented by a two-sheeted hyperboloid in three-dimensional Minkowski space.

Since hyperbolic geometry expands exponentially, it enables us to create more space in a limited environment than we could in Euclidean geometry. Unfortunately, this does come with issues for computers. Since the space expands exponentially, we would need to have a high numerical precision to be able to calculate the positions of objects in the space. To get around this, we will be using a tree structure that defines the location of objects instead of their actual hyperbolic coordinates. When needed, we can translate the positions in the tree to their relative hyperbolic coordinates. This type of non-Euclidean geometry is used in Holonomy VR to create an infinite world in a limited physical space.

2.2. Holonomy VR

Holonomy VR [50] is an application that uses hyperbolic geometry and the concept of parallel transport to create an experience in which is it possible to explore an infinite world in a limited physical space where the only form of locomotion is walking. This is accomplished by creating a virtual environment (VE) that lets users explore the hyperbolic plane in VR. In Holonomy VR the VE is built up of tiles that are square. In VR these tiles seem to be laid out in a 3 by 3 grid and are surrounded by bushes where the user cannot walk through. This means that the user can never walk more than 3 steps in the same direction and that they are not able to see tiles outside of the grid. However, since we are using hyperbolic geometry we can actually fit more than nine tiles into this grid. On a Euclidean plane we can fit exactly four squares around a single point. However by creating a closed loop of five squares around a single point instead of the normal four squares we create a hyperbolic plane. This form of the hyperbolic plane is called the order-5 square tiling, see Figure 2.3. Since there are now five squares around a single point instead of the usual four, the physical location does not correspond the virtual location when we move around a point. As explained previously, moving in a loop through non-Euclidean space will cause the direction to change. This is called holonomy and is where the application gets its name from. With this principle we can fit at least 13 up to 17 squares in the 3 by 3 grid depending on the user's location.

With this setup we can create our own hyperbolic world which users can explore. Currently Holonomy VR features two modes, finding the flag and collecting keys to open a chest. Both of these modes let the user explore the hyperbolic plane and require them to use the properties hyperbolic space to reach their goal. To help them navigate, the user is given access to a map that shows them a heatmap based on the objectives that they need to reach.

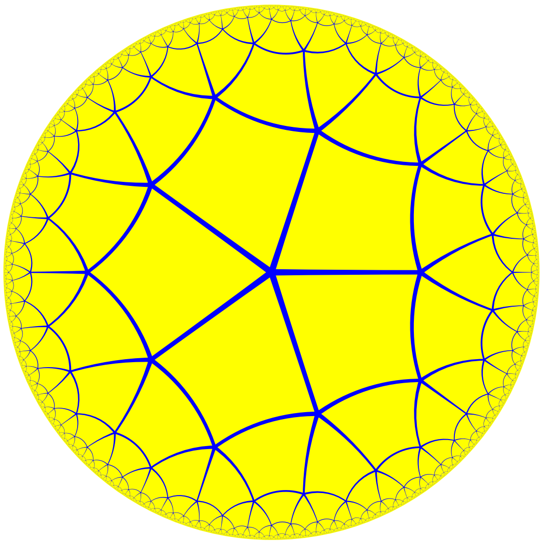

Figure 2.3: Order-5 Square Tiling [53]

3
Related Work

This section explores the related work to this thesis. We start by learning about navigation and spatial cognition. Then we take a look at the adaptability of humans which supplies us with the theory that we could navigate in hyperbolic space. Next, we explore the related work about non-Euclidean geometry and finally, we look at the concept of affordance and how it relates to the properties of an environment. This section will provide us with the necessary background information to understand how we navigate in Euclidean space and how we can make this research applicable to non-Euclidean environments.

3.1. Euclidean navigation

We start by exploring the related work about navigation. It is an important factor in our everyday life. We need a form of it to get from point A to point B. Naturally, the default space in which navigation is researched is Euclidean, since the physical world that directly surrounds us behaves as such.

Navigating in Euclidean scenes has been explored quite extensively. There is related work about navigating in virtual versus physical environments, how to train your navigational ability for specific environments, and how this training can help create your cognitive map. Cues are also an important factor in navigation and can help users to navigate more efficiently. The navigational performance can also be influenced by the characteristics of the user. This is important to take into account when designing a user study. But first, we start with navigating in virtual and physical environments.

3.1.1. Virtual versus physical

An important variable in our research is virtual reality. It allows us to create environments that are not possible in the real world. To see if the research about navigation is applicable to our research, we need to know if there is a difference in navigating in VR compared to navigating in the real world. Since VR is a way to experience VEs, we will start by looking at established research that compares navigating in a VE to navigating in the real world.

Koenig et al. [30] have assessed navigation in real and virtual environments. They used multiple screens to cover nearly the entire view of the participant, which made the experience more immersive. They found that there is no significant difference in navigating real and virtual environments. However, some participants had difficulty controlling their movements through the VE due to a lack of experience with the controls, which reduced their navigational performance. Thus, the performance of users does depend on their experience with controlling the VE.

Next, Marin-Morales et al. [32] compared navigation between VR and real life. They have established that there is no significant difference in navigating real life and in VR. The participants that wore the VR headset could move through the environment by walking within a limited space and using teleportation. They argue that this result relies on the immersive aspect that a VR headset gives, which implies that if a user does not wear a VR headset, they feel less immersed and thus perform worse.

Combining these ideas, Hejtmanek et al. [23] compared navigation in the same environment while being seated, wearing a VR headset, and using the actual real-life environment. They conclude that participants need more time in VEs to achieve the same level of spatial knowledge. They argue that this is caused by reduced immersion in the VE. For example, there were posters present in real life, but not in the VE, and participants indicated that they relied on those posters. Another aspect that influences spatial knowledge is again the controlling of movement. They used an omnidirectional treadmill for movement with the VR headset; however, omnidirectional treadmills do not engage the full range of somatosensory receptors and muscles compared to stepping in the real world which could have influenced their results.

We can conclude from these previous works that the navigational performance of humans navigating in VEs can be comparable to navigational performance in real life, with a side note that we need to be careful of a few aspects, such as how the participants are able to move through the VE, how the VE compares to the real environment, and how the environment is presented to the user. This comparison allows us to dive deeper into the research about navigation.

3.1.2. Training

Before humans can navigate efficiently, they need to learn about an environment. In practice, this training is done by looking at a map, following a virtual tour, or by actively exploring a virtual or real environment. The next researchers discuss the benefits of these learning methods.

Richardson et al. [46] looked at learning by actively exploring a VE and by studying a map. The exploration in the VE is done while participants are seated. They found that, with an equal amount of exposure, the spatial knowledge of the participants was equal compared to learning by exploring the real environment. However, when the map had a misaligned orientation, the performance of participants plummeted.

Farrell et al. [16] had the same idea, but they also included learning by exploring a VE with a map, by watching the researcher complete the training, and by exploring the VE using a line that guided them straight to the targets. They also included a control group by not letting participants train. They found that training by actively exploring the VE (with or without a map) and by only using a map produced better navigational results compared to no training. The results of these methods were also equal. However, the training times differed. Only using a map was the fastest, then exploring while having access to the map, and at last exploring without access to the map. They argue that the performance of these training methods is the same because participants did not use the egocentric information that was available in the VE. This is probably due to the fact that the VE that they used was not immersive. The passive experiences did not improve navigational performance after training. Most likely, this is the effect of passive learning since the participants were not actively exploring the environment.

These were both fairly old works and VEs have become more immersive and realistic. A newer study by Snopková et al. [51] compared an active virtual tour, i.e., being able to look around and click in a direction to go back or forward in the tour, and a map training. The tour was again being done while seated and was created with 360° photos. They conclude that both methods produce equal navigational performance when performing the actual task in the real environment. However, participants that followed the tour were able to provide more detailed navigational instructions and include more landmarks and visual characteristics in these instructions than the other participants. This is most likely due to the fact that the tour included more visual details.

Active training with VEs and maps are thus a great way to improve navigational efficiency. However, the completeness of the training environment is very important. Missing visual information can result in degraded performance. Another thing to note is that training time when only a map is used is shorter, but comes with the downside that the orientational performance will be lacking. Combining active training in VEs together with using a map seems to be the way to go when we want the lowest training time while also achieving adequate orientational performance. We have now seen that training can improve navigational performance, but why does it improve? That has to do with cognitive mapping which is discussed in the next section.

3.1.3. Cognitive Mapping

Cognitive maps are created when we explore or navigate in an environment. The cognitive map can be created by learning about landmarks in an environment. We remember the route in the cognitive map and can figure out shortcuts from other routes. It allows us to approximate the distances and angles to different waypoints and make decisions based on these memories. We can also navigate without landmarks by the form of dead reckoning, but the accuracy reduces the further we travel [39].

There is extensive related work about cognitive mapping in humans but also in animals. The cognitive map can be tested by asking questions about landmarks and routes but also by creating drawings. However, these questions need to be picked carefully to be able to evaluate the cognitive map. The next researchers discuss different methods to evaluate the cognitive map.

Kitchin [29] lists multiple methods to test the spatial knowledge of participants. The study was carried out with subjects that were students who lived in the used area for about 10 weeks and they had lessons on maps and had tours around the area. He used multiple methods to test their cognitive map:

- Drawing a map
- Completing a partial map
- Listing distances between landmarks
- Listing angles between landmarks
- From three landmarks, underline the landmark that is the furthest from the other two landmarks
- From a given place, draw a line to a given landmark with the correct orientation and length
- Validate a map

Another interesting study is done by Epstein et al. [14]. They compared the research about cognitive mapping in rodents with that of humans. They found that the results of the research about rodents match that of humans.

Most of these methods seem difficult to apply to hyperbolic geometry, especially methods that have to do with maps. However, the methods applicable to landmarks could be used to evaluate the cognitive map in hyperbolic geometry. How effectively the cognitive map is created is very much based on which cues are used. This, we will discuss in the next section.

3.1.4. Cues

There is established research about the effects of navigational cues on navigational performance and spatial knowledge. The latter is more important when we want to know how humans effectively learn about the spatial layout of an environment [28]. Khan et al. concluded that using a map is an effective cue for navigation but hinders users from understanding the spatial layout of the environment. It would instead be better to promote active exploration of the environment by giving them verbal instructions. Then participants will still be able to efficiently navigate, but also be able to recall landmarks and indicate angles and distances between landmarks more effectively. Parush et al. [37] evaluated if there is a difference in using automatic navigation systems, i.e., always knowing your location, versus actively asking about your location. The automatic navigation system always keeps you updated, which does improve navigational performance. However, Parush found that requiring the subject to actively ask about their location increased their spatial knowledge significantly.

Navigational performance can be aided by navigational cues. Here Parush et al. [38] compared a map to a route list, i.e., a list of directions, in an environment with and without landmarks. They found that using a map took more time, but in the end was as effective as the route list. Moreover, performance degradation upon removal of the navigational cues was less for those that navigated with landmarks as compared to no landmarks, indicating that landmarks are important for navigation.

Next, Burigat et al. [6] compared three different navigational cues: a 2D and a 3D arrow pointing to the location of the target and a radar. It was found that the 3D arrow outperformed the other cues in terms of navigational performance. Chan et al. [8] discuss the function of visual location information in spatial navigation. Visual location information can be described as a beacon. Beacons provide accurate positional information and are, by definition, highly reliable indicators of a specific location.

Finding beacons even emerges early in the development of humans and it appears to be one of the most basic forms of landmark-based navigation.

These researchers all use different approaches and methods to execute and analyze their results. To make comparisons between studies easier, Wiener et al. [60] presented a taxonomy of wayfinding tasks that distinguish them based on external constraints as well as the information available to the user.

Using cues is a good way to make sure that people can start efficiently navigating an environment without any knowledge about the environment. It might take them more time than trained users depending on the cue, but they do not need to take the training and it does result in adequate navigational performance. A factor that can also influence navigational performance is the characteristics of the user. This is discussed in the next section.

3.1.5. Characteristics

Human characteristics can influence navigational performance. For example, Wiener et al. [61] found that age differences can influence navigational performance. They found that older adults performed worse in the route-repetition task, the route-retracing task, and the directional-approach task. They argue that this is due to the fact that older adults have a reduced ability to learn new information and to remember it. This indicates that age differences can influence navigational performance.

Another characteristic that can influence navigational performance is gender. Munion et al. [35] found that men perform better than women. They followed persons in the real world using GPS trackers and found that ”men produced more directional persistence, less pausing, and less revisiting compared with women, which in turn related to more navigational success”.

These characteristics are important to take into account when designing our study about navigational performance. It is important to have a balanced or very specific user study group to counteract the differences in navigational performance due to these characteristics. Much information is thus available about navigation in Euclidean scenes. But before we continue to apply these methods to hyperbolic geometry, we need to research if this knowledge could carry over to hyperbolic geometry. We will do this by looking at the general adaptability of humans.

3.2. Adaptation

From previous work and experience, we know that humans are very adaptable to changes in their environment and alterations in their perception. For example, Bock et al. [4] experimented with rotated vision. They let participants point to an object while the participants cannot see their hands directly. Instead, a pointer is projected onto a display that tracks the location of their hand, but this view is rotated. They found that humans adapt quite fast to the rotated vision; however, participants had trouble adapting to a rotation of 90 degrees.

A great source to find previous work about adaptation is the book ”Perceptual Modification: Adapting to Altered Sensory Environments” by Robert B. Welch [59]. Although the book is quite old, it reviews established work about adaptation from the previous decades, which is still relevant today. In the book, adaptation is defined as something that only exists when someone experiences a negative aftereffect when the interference is removed. This definition is a more extreme version than intended in this thesis, since we do not expect people to experience a negative aftereffect.

The amount of adaptation can also change based on whether the experiment is active or passive, i.e., if the subject needs to actively move, even their eyes, or if the subject is sitting still. The ”Reafference Hypothesis” hypothesises that this fact is always true, but there are many exceptions; however, the amount of adaptation when active is often higher than when passive. There is most likely also a difference between learning and adapting; this is due to the fact that learning does not produce negative aftereffects, while adaptation does.

The book includes many studies about adaptation, some more successful than others. But, Welch concludes that, from the studies discussed, humans are (partially) adaptable to: prismatic displacement, vision transposition, optical tilt, placement of rearranged audio, and distortions in visual stability, depth, distance, and form. This adaptation is rarely influenced by the characteristics of the subjects; if they are healthy, they show roughly the same amount of adaptation. The amount of adaptation does differ per interference, just like the adaptation time. Some adaptations only take minutes while some can take days. An important thing to note is that adaptation only occurs when there is a motivation for it, being either by rewards or punishments. These do not have to be large; for example, pointing in the wrong direction is already unpleasant enough. Figure 3.1 shows a flow chart for the proposed general model of adaptation by the book.

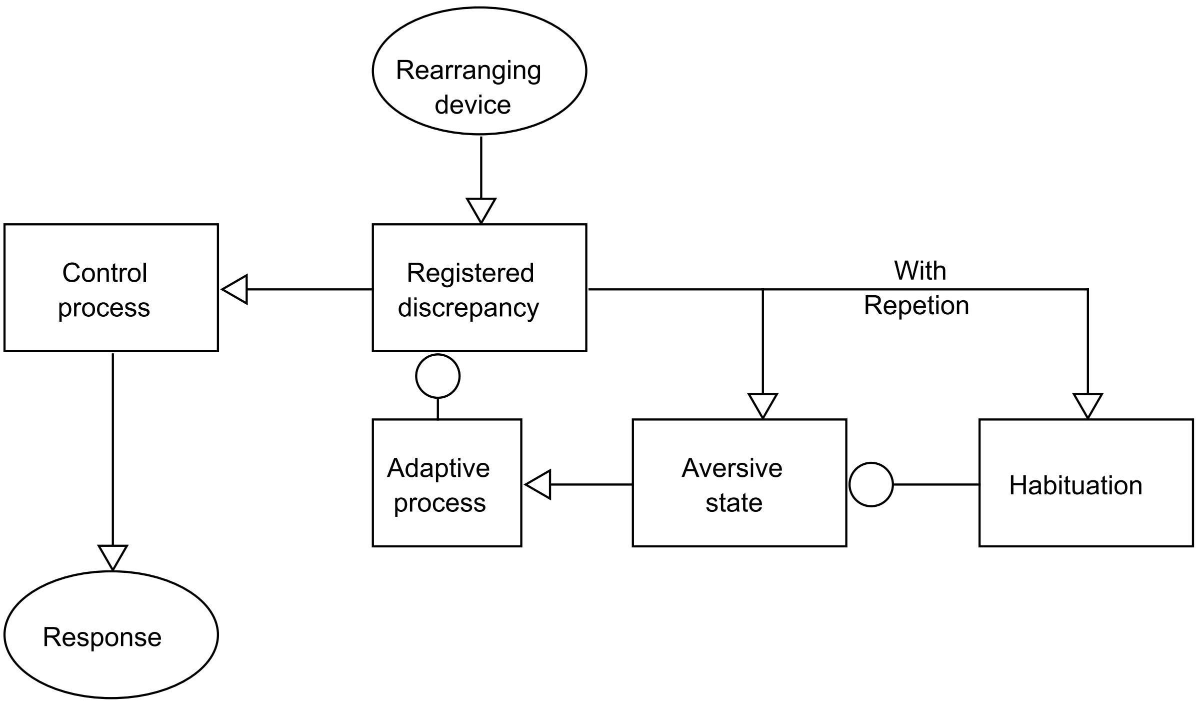

Figure 3.1: Proposed general model of adaptation [59]

Although this work is not specific to navigating in environments, it does show that humans are flexible and adaptable in their thinking and behaviour, which indicates that humans could be able to learn how to navigate in a confined hyperbolic space, especially when training specifically for it. Next we will take a look at research on how to create non-Euclidean environments, how people experience them and how they might navigate in them.

3.3. Non-Euclidean Navigation

Non-Euclidean ways of navigation can be achieved using different methods. First, we take a look at redirected walking, then we explore the use of portals and at last we take a look at navigation in actual hyperbolic geometry. All of these methods are used to create non-Euclidean geometries in virtual environments and allow users to explore and interact with spaces that are different from the familiar Euclidean geometry.

3.3.1. Redirected Walking

An easy way to create a non-Euclidean environment is by redirected walking. This is a technique used to allow users to walk through large-scale VEs while physically remaining in a reasonably small workspace. This is achieved by intentionally injecting scene motion into the VE by rotating the user's virtual orientation without them realising it, which can imperceptibly alter the user's path. This creates the illusion of walking in a straight line while actually following a curved path. The technique can be used to create more immersive and engaging virtual experiences, as users can explore VEs more freely and naturally since the limitation of physical constraints is reduced. When done well, the user will still think that they are in a Euclidean environment and will not notice the redirection.

In 2001, Razzaque et al. [44] studied how much a VE can be altered in terms of rotation and distance before it becomes noticeable. A year later [45] they used this idea in a square room setup that projects a screen onto three different walls while leaving one blank or open. Their goal was to make the blank or open wall less noticeable by exploring the use of redirected walking. They let users walk in place while the VE rotates around them. They found that most users did not notice the rotation of the VE and that it reduced the need for a controller to rotate the environment on their own. Later, this idea was also explored by Steinicke et al. [52].

3.3.2. Portals

Portals can be used to create non-Euclidean geometries. Portals enable users to traverse between different locations instantaneously, creating the illusion of impossible spaces and altering the perception of the virtual environment. This can be used to create mazes, puzzles, and other interactive experiences that challenge the spatial understanding and navigation skills. However, portals can also be used to seamlessly connect different areas of a virtual environment, allowing users to explore large-scale spaces without the need for complex navigation or locomotion techniques. Next, we take a look at the actual research in navigation that uses portals to create non-Euclidean environments.

Warren et al. [57] used wormholes in a maze and observe if people take these wormholes after exploring the maze to find an objective. They research if humans create a Euclidean map in their mind with metric distances or a cognitive map with landmarks and paths between them, or a combination. The conclusion about the cognitive map is unclear. However, they did observe that people are equally good at navigating the Euclidean maze as the non-Euclidean maze.

Next to wormholes, more complicated impossible geometries are also explored. For example, Mayor et al. [34] explored the impact of altered geometries on the user's sense of presence. The findings suggest that while non-Euclidean elements can initially disorient users, they can also enhance the overall immersive experience by providing novel and engaging interactions that are not possible in traditional Euclidean spaces. The article concludes that the use of portals significantly contributes to the richness and depth of VR experiences and allows for more movement than the actual boundaries of the physical space would allow.

Portals are thus a powerful tool to create non-Euclidean geometries in virtual environments. They allow the creation of seamlessly transitioning spaces. This method is also incorporated in Holonomy VR to let users explore the hyperbolic plane in VR. But the question remains, can people actually navigate in hyperbolic space? This is discussed in the next section.

3.3.3. Hyperbolic Geometry

Hyperbolic geometry has been used in virtual environments to create immersive experiences that simulate the properties of hyperbolic space. For example, hyperbolic geometry has been used to create virtual environments in games like HyperRogue [31] and Hyperbolica [10]. These allow users to explore and interact with hyperbolic surfaces and spaces and create novel and engaging experiences that challenge the user's spatial understanding and navigation skills.

Previous work about exploring hyperbolic geometry is scarce. Hard et al. have explored the viewing of hyperbolic space H^3 [19] and viewing the two-dimensional hyperbolic space combined with one-dimensional euclidean space $H^2 \times E$ [20] in VR. They created a program that lets users walk around in a small space of this environment and experience the effect of holonomy that hyperbolic geometry creates. Next to this, there are games that let you explore hyperbolic geometry. For example, the 2D game HyperRogue [31], where the authors found that backtracking to a location already visited is found difficult by the users. Another example is Hyperbolica [10], which creates a 3D hyperbolic experience where the user needs to complete tasks in hyperbolic space and use it's properties to their advantage.

There is also research about navigating on the hyperbolic plane in VR [41]. By following balls placed on a hyperbolic plane, Pisani et al. found that humans are able to find the balls efficiently and in the correct order. However, this task is quite straightforward since the users were always able to see the balls and approach them directly. The task of navigating in a hyperbolic space can be complex and probably requires the user to understand the actual properties of the hyperbolic geometry. It might be possible to let users learn about these properties by letting them actively interact with the hyperbolic space. They could then learn them through the concept of affordance.

3.4. Affordance

The concept of affordance is defined by Gibson. In his book "The Ecological Approach to Visual Perception" [17], Gibson describes in great detail how we, humans, observe the environment around us. The brain sees the geometry of the environment and the objects within it and relates that visual knowledge to possible actions. As such, an environment affords actions that we can take to manipulate objects in the environment, but also to move through the environment. In "Affordance for Spatial Navigating" [18], Gregorians reviewed the history of established research about affordance and spatial navigation. What an environment affords is a key aspect of how we navigate the environment. Walls, boundaries and landmarks are key examples of affordances that allow us to evaluate where we can and cannot go.

What hyperbolic geometry affords has not been researched. But we can theorise about its affordances. Hyperbolic geometry has some unnatural properties compared to Euclidean geometry. In this way, hyperbolic geometry can afford the user actions that are not possible in a Euclidean scene. For example, being able to use holonomy to rotate the environment. However, the user must notice these possible actions to be able to utilise them in a way that is beneficial for the user. If they do not notice these actions, they might not be able to efficiently navigate and relate these actions back to the properties of hyperbolic geometry. Fortunately, as discussed previously, we do know that humans can navigate in hyperbolic geometry, which makes it also likely that they do notice what hyperbolic geometry affords and relate these affordances back to the properties of hyperbolic geometry.

4 Main Contributions

Holonomy VR lets users explore hyperbolic space in virtual reality and with a map. But can people be trained to efficiently navigate in an unnatural form of geometry, like hyperbolic geometry, and also learn about this geometry? From research about navigation in Euclidean scenes, we know that humans are able to learn to navigate an environment efficiently using different training methods. We also know that humans can explore the hyperbolic plane in VR. Since humans can adapt to new experiences they should be able to learn to navigate effectively in hyperbolic space. But which training method would perform the best, or will they result in the same performance? Another question that arises is, do people learn the properties of hyperbolic geometry when they are learning about navigation in a hyperbolic scene?

By answering these questions we can verify if training methods for navigation in a Euclidean scene can be used for training for navigating in a hyperbolic scene. Using the different training methods we can explore if a method is more functional than another in terms of exploring a hyperbolic scene but also to learn about its hyperbolic properties. Unlike with Euclidean geometry, there is a reasonable chance that these differences exist. First, let's look at the VR environment in Holonomy VR. In this environment, users might get a better feeling for what their actions will imply. Before they take a step, they can see how the environment will change without taking the actual step. By thinking about what their action will imply, users might need fewer steps to actually reach their goal. With a map, this same information is available but is not interactively promoted to the user.

Using a map also has its benefits. Users using the map will get direct feedback from their surroundings when they move to another square. They will see that inaccessible tiles are now accessible and they will see that the environment has been rotated, see Figure 4.1. This property of the map might indicate that users using the map view will have better impressions of the properties of hyperbolic geometry. With this feedback, users might also be able to navigate more efficiently. Another benefit of the map is the fact that users can see further, and that they might get a better overview of the scene.

Assuming that research about navigation in Euclidean scenes can be applied to hyperbolic scenes, we would expect that the training methods give adequate navigational performance. However, there might be a difference in learning about the properties of the hyperbolic environment, which might play a role in people's navigational performance. If this is the case, then some training methods might be better suited for hyperbolic geometry than others.

4.1. Research Objective

The aim of this thesis is to investigate if humans can navigate in hyperbolic space when their movement range is confined and to see if they learn about the properties of hyperbolic space, i.e. the curvature and the exponential growth of the space. This is done by measuring their navigational performance, which we define as how efficiently someone navigates to a target, meaning the fewer hyperbolic tiles they visit, the better. However, we cannot expect users to directly be able to navigate in an unknown hyperbolic space. Just like in Euclidean geometry, they need to familiarise themselves with the environment before

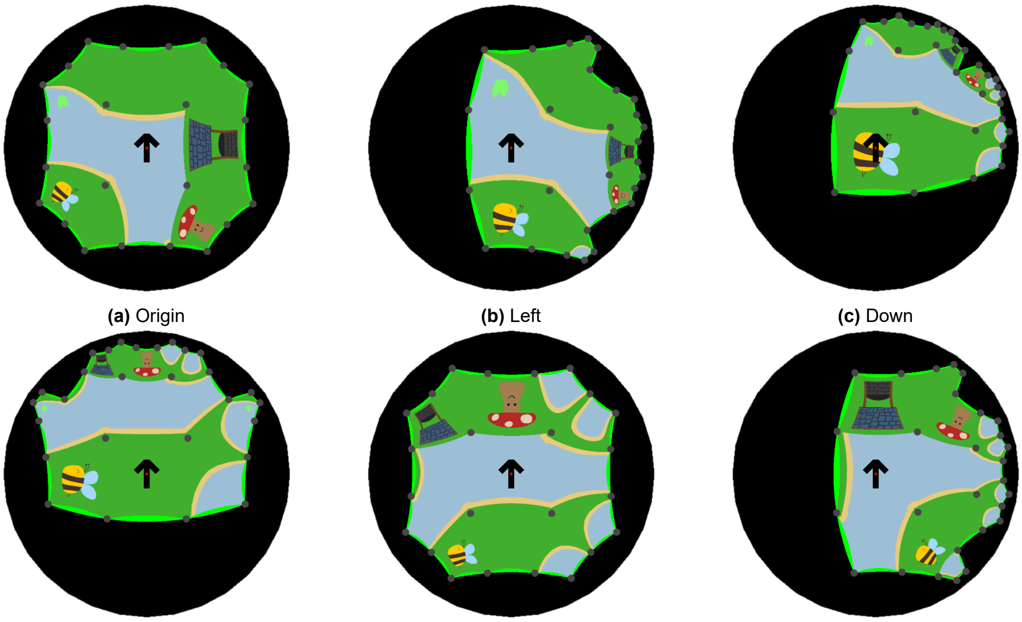

(a) Origin **(b)** Left **(c)** Down **(d)** Right **(e)** Forward **(f)** Left

Figure 4.1: Traversing around a single point on the hyperbolic plane reveals inaccessible tiles and causes the plane to rotate

they can navigate efficiently, or they need to use navigational cues to help them find their way to the target. Since we want to encourage learning about the environment, we will focus on familiarisation. Familiarisation can be achieved through training. The following training methods have shown to be effective: exploring a real-life scene, exploring a virtual scene, or using a map. The map can be used standalone or in combination with exploring. Training is discussed extensively in subsection 3.1.2.

The training methods mentioned produce equal navigational performance if users are familiar with the controls. However, the methods do take different amounts of time to complete, which defines the training time. We will use these methods to train people to navigate in a confined hyperbolic space. We assume that this is possible since humans are adaptable and can learn to explore hyperbolic space when their movement is not limited [41]. We can then explore if results from navigational training in Euclidean space carry over to hyperbolic space. If this is not the case then this might be related to understanding the affordance of the geometry [17]. The affordance boils down to the actual properties that the hyperbolic space presents to the user.

If this study is successful we can use Holonomy VR to teach people about hyperbolic geometry. It could also be used in VR games to allow for novel game mechanics and to provide a natural way of locomotion, namely walking. This natural way of locomotion has the benefit that there is no illusion of motion that users experience, which is one of the main causes of VR motion sickness [58, 36]. This leads us to the following research questions.

> *Main RQ: How does navigational training in a hyperbolic environment influence navigational performance in this environment and the understanding of its hyperbolic properties?*

To answer the main research question, we can split it into several sub-questions that address different aspects of the problem. These sub-questions will help us understand how learning about a confined hyperbolic environment influences navigational performance and the understanding of its hyperbolic properties.

> *RQ1: Can humans adapt to be able to explore a confined hyperbolic space?*
>
> *RQ2: Which established training method has the shortest training time for navigating in hyperbolic space?*
>
> *RQ3: Which established training method produces the best navigational performance in hyperbolic space?*

RQ4: Are the results from navigational training in hyperbolic space comparable to those in Euclidean space?

RQ5: How well do people understand the properties of hyperbolic space after training with different methods?

RQ6: Does the understanding of hyperbolic properties improve navigational performance in hyperbolic space?

4.2. Hypotheses

For the first question, we know that humans are quite adaptable and we know that humans can explore hyperbolic geometry. The next three questions can be answered if we assume that previous work about navigating in a Euclidean environment carries over to a hyperbolic environment. The last two questions are related to the way an environment is experienced.

We hypothesise that for *RQ1* humans are able to adapt. We know that humans are already able to explore hyperbolic space when their movement is not limited [41]. We also know from previous work that humans are adaptable and are thus likely to be able to adapt their thinking to be able to explore and navigate hyperbolic space when their movement range is limited.

This is related to the change process. This process is visualised in Figure 4.2. When we confront people with our environment we introduce them to hyperbolic geometry. In our case, this will not cause resistance since people should be willing to learn about the environment. However, they will end up in chaos since they do not know how to explore the environment efficiently and will get easily lost. But we assume that they will learn from this experience and get a transforming idea which they can integrate into their ability to explore the environment. This will then lead to an understanding of the environment and they will be able to explore a confined hyperbolic space.

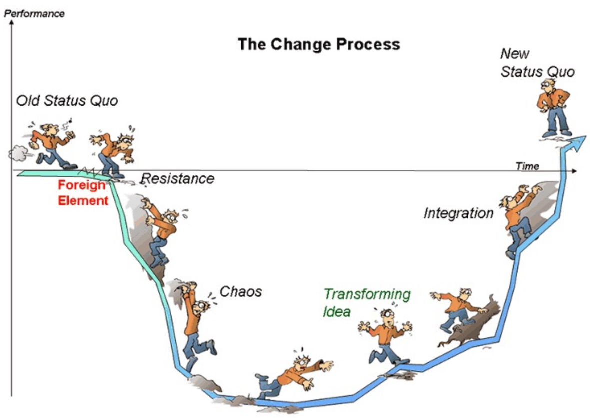

Figure 4.2: The change process that people go through when they adapt to the environment [15]

Hypothesis 1 (H1): *Yes, humans can adapt to explore a confined hyperbolic space.*

For the training methods, we first hypothesise that for *RQ2* training with a map would be the fastest, then training with the environment and map, and finally only training with the environment. This is based on established research that shows that completing navigational training can take a different amount of time depending on the training [16].
Hypothesis 2 (H2): *Training with a map would be the fastest, then training with the environment and map, and finally only training with the environment.*

Then we hypothesise that for *RQ3* we expect from established research that people have roughly the same navigational performance after following training [46, 16, 51].
Hypothesis 3 (H3): *Navigational performance after following training is roughly the same between different training methods.*

For *RQ4* we hypothesise that the results from navigational training in hyperbolic space are comparable

to those in Euclidean space. This is based on the assumption that humans are adaptable and that they can adapt to explore hyperbolic space.
Hypothesis 4 (H4): *The results from navigational training in hyperbolic space are comparable to those in Euclidean space.*

For *RQ5* we hypothesise that the navigational performance after training is adequate, and thus, in each training, we learn the affordances that the hyperbolic environment presents. This is based on the previous work that concludes that what an environment affords is important for spatial navigation, it determines where we can and cannot go [18]. If we are not aware of the affordances of an environment we cannot efficiently navigate. From the third hypothesis, it is expected that the navigational performance after training is adequate, and thus, in each training, we learn the affordances that the hyperbolic environment presents. So being aware of the affordances we should be able to relate it to the properties of the hyperbolic environment.

In Holonomy VR the affordances of the hyperbolic environment differ between the actual VR experience and the map experience. In the VR experience, the affordances in each tile correspond to the same affordances in Euclidean space since a tile is Euclidean. However, when looking at multiple tiles the user can see that tiles behind trees differ depending on which side of the tree they look. This gives the user a choice before they take the step. This is different from the map experience where the user can only see the environment change after they take a step. However, the map does allow them to see how many tiles there are around a single point, which tiles are available outside of the border, which new tiles are now accessible when they take a step and that the environment rotates when they walk around a tree. Both experiences thus give different affordances, however, the affordances still relate to the hyperbolic properties.
Hypothesis 5 (H5): *In each training, we learn the affordances that the hyperbolic environment presents. Being aware of the affordances we are able to relate it to the properties of the hyperbolic environment.*

For *RQ6* we hypothesise that the understanding of hyperbolic properties improves navigational performance in hyperbolic space. This is based on the assumption that the understanding of the affordances of the hyperbolic environment is important for spatial navigation.
Hypothesis 6 (H6): *The understanding of hyperbolic properties improves navigational performance in hyperbolic space.*

The hypotheses will be validated by performing a user study. The user study will consist of a practice, training and evaluation phase. The practice phase allows participants to get familiar with Holonomy VR. The training phase will consist of the participants learning about the landmarks contained in the hyperbolic space by exploring it or by looking at a map. And the evaluation phase will consist of the participants navigating in the hyperbolic space to specific landmarks.

The first and second hypotheses will be measured by the time taken to complete the training. The third hypothesis will be measured by the number of tiles visited to complete the task after the training. The fourth hypothesis will be validated by looking at the results of the second and third hypotheses and comparing them to the results found in Euclidean space. The fifth hypothesis will be validated based on the number of correct answers to multiple-choice questions about the properties of hyperbolic space. Finally, the sixth hypothesis will be validated based on the fifth hypothesis and the users' final navigational performance.

Using the hypotheses above, we hypothesise that navigational training in a hyperbolic environment improves navigational performance since the research about navigation in Euclidean space will be comparable to hyperbolic space, and thus training will improve the navigational performance. We also hypothesise that people will learn about the properties of hyperbolic space when they are training for navigation in a hyperbolic scene. This is based on the assumption that the affordances of the hyperbolic environment are important for spatial navigation, and that the understanding of the affordances will improve navigational performance.
Main Hypothesis: *Navigational training in a hyperbolic environment will improve navigational performance in this environment and the understanding of its hyperbolic properties.*

5
Technical Contributions

Before we can evaluate the hypotheses, Holonomy VR must be updated to support the technical requirements. The map is updated to improve the performance and allow for the level of detail to experience Holonomy VR in 2D. The beacon cue is added to guide the user to the objectives without the need for a map. The landmarks are added to create a new game mode where the user needs to find landmarks in a specific or non-specific order which will be used as the navigational task in the user study. The user study menu is added to allow for our user study to be conducted. The results of the user study are saved and can be analysed to evaluate the hypotheses. The improvements are discussed in the following sections.

5.1. The Map

The previously provided map in Holonomy VR gave a good basis but it was limited. The map did not show movement in the current tile and it used discrete transitions. These limitations existed because of performance reasons [62]. To keep Holonomy VR running well, the depth of the map was limited to a depth of 3 tiles. The performance issue was caused by the fact that the map was sent to the GPU on a per tile basis. This is not an issue if the map is only updated after transitioning to another tile, however, we quickly run into frame rate issues if we want to update the map every frame. To increase the performance, we can send all the tiles of the map to the GPU at once, this reduces the CPU overhead substantially.

Unfortunately, performance was not the only issue. The drawing of the map was not correct since the edges of the tiles were overlapping. Another problem which was revealed during texture mapping is that the tiles did not keep a consistent rotation. After solving these issues, the map is extended with textures and allows for updating with the movement of the user. The performance of the indicator is also improved and finally, the map can now be used to explore Holonomy VR in 2D.

5.1.1. Performance

To increase the performance of the map, we need to send all the data from the map to the GPU in one go. We do this by using the TextureData class. This class stores all the information of the map in a format that can be directly sent to the GPU. On the GPU, we use a compute shader that checks in which tile the pixel lies, and we update the colour of this pixel based on the tile data. With this improvement, we can now render the map with a depth of 4. However, the depth is now limited by the maximum vector size of the compute shader in Unity. In theory, we could work around this, but it is expected that it will take a hit on performance since the number of tiles on the map grows exponentially with the depth. Now that we have the performance that we need, we can correct the drawing of the map.

5.1.2. Geodesic Edges

On the map in Holonomy VR the hyperbolic space is mapped to the Poincaré disk. For the corners of our tile we can use a simple mapping from the hyperboloid model used by the spanning tree to the Poincaré disk, but then how do we display the edges? Previously the edges of the tiles were overlapping, as can be seen in Figure 5.1. This happened because a rough check was performed to see if a pixel lay in a certain tile. However, this check did not result in a unique tile location. So, how should we calculate the correct edge?

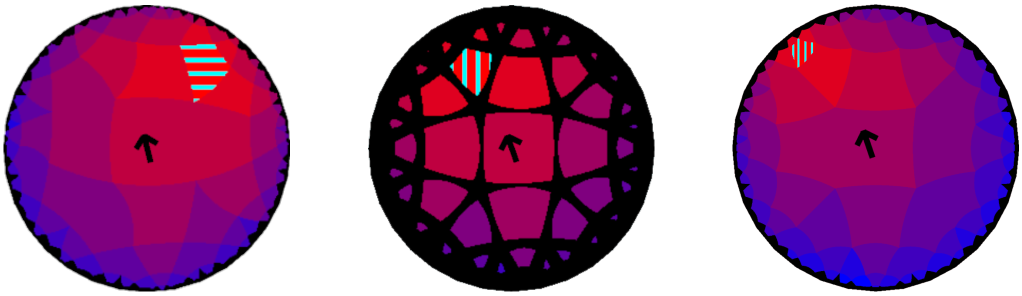

(a) Tiles are overlapping, which results in an inconsistent tiling

(b) The black lines between the tiles are the amount of overlap

(c) Example of the correct map

Figure 5.1: Example of the incorrect tiles

In Euclidean space the edge between two corners is the shortest line that we can draw, which is just a straight line. However, when mapping the hyperbolic plane to the Poincaré disk this line is actually a curve, except when we draw a line that intersects with the centre of the disk. This line is called the geodesic line, which is also called the limiting parallel [21]. This line intersects the unit disk at its ideal points and gives us the shortest path between our two points. The ideal points lie on the unit disk and thus lie on the end of the hyperbolic plane, so at infinity. But how do we calculate this line? First of all we need to find its ideal points. Given two points A and B, their ideal points are the intersection points of their orthogonal circle with respect to the unit circle [56]. We find this circle by inverting one of the points with respect to the unit circle. This gets us three points: A, B and the inverse of A which we call A'. These points define the orthogonal circle, if the three points are on one line, then they describe a straight line, which indicates that the ideal line crosses through the origin. The two points where the unit circle and the orthogonal circle intersect are the ideal points and the part of the circle between the points A and B is the limiting parallel, see Figure 5.2a. This gives us our definition for an edge on the map, see Figure 5.2b to see the lines on the map.

The above is a complex way to calculate the geodesic edges and is not used by Holonomy VR. Instead we compare the hyperbolic distances of the centres of the overlapping tiles and use this to calculate the geodesic edge. The calculation of the ideal points can be used to calculate the hyperbolic distance, but this is not very efficient. Instead we calculate the distance in the hyperboloid model [43]. This is done by taking the dot product of the two points, which is the Lorentz inner product in the hyperboloid model. Then we take the inverse hyperbolic cosine of this negative value:

$$d(\mathbf{u}, \mathbf{v}) = \cosh^{-1}(-\langle \mathbf{u}, \mathbf{v} \rangle_L)$$

where $\langle \mathbf{u}, \mathbf{v} \rangle_L$ denotes the Lorentz inner product of the points $\mathbf{u}$ and $\mathbf{v}$. This gives us the hyperbolic distance between the two points, which we use to approximate the hyperbolic edge. Now that the map looks correct, see Figure 5.1c, we can move on to the next step, moving the map with the user.

5.1.3. Moving with the user

Moving the map with the user is done by using a subset of Möbius transformations in the Poincaré model. These transformations keep the hyperbolic distance and orientation between points the same when transforming them with the same Möbius transformation. A general Möbius transformation is calculated by $f(z) = (az + b)/(cz + d)$. The coefficients a, b, c and d for a specific transformation can be found by creating a mapping of three points to three other points. In the Poincaré model we can calculate the Möbius transformation between two points a and b by calculating their ideal points and mapping the ideal points to the same place and mapping a to b. This causes the ideal points to always stay in the same place. Since we always map two of the three points to the same place we only use a subset of the Möbius transformations.

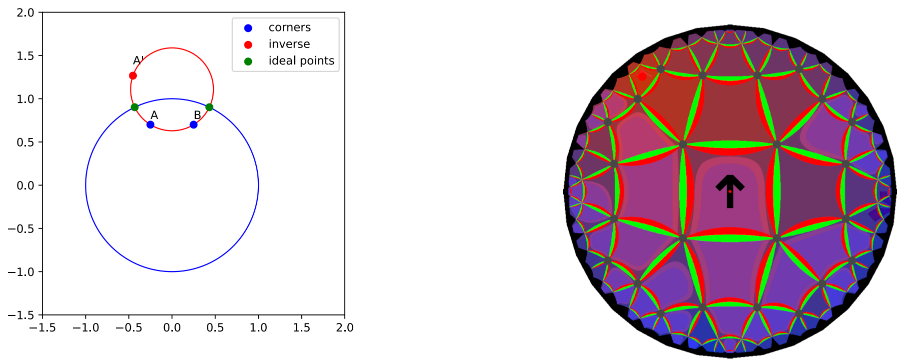

(a) The limiting parallel and its ideal points

(b) The map with the limiting parallels. The limiting parallel is the line between the red and the green.

Figure 5.2: The limiting parallel in practice

When the map is generated the centre of the centre tile is always at $0, 0$. So we only need to calculate the hyperbolic position of the user. This calculation is quite simple because we want to know the distance from the user to the origin of the map. Which is the inverse of the hyperbolic tangent of the Euclidean distance d of the user to the origin: $\tanh^{-1}(d)$. We also need to know the hyperbolic angle from the origin to the user. Again, since we use the origin the hyperbolic angle is equal to the Euclidean angle. This angle keeps the rotation of the user in the coordinate. Then the only thing left is to calculate the Möbius transformation from the centre to the user position and transform every tile with this transformation.

However the indicator must also dependent on the user position. For that we store the hyperbolic x and y distance of the user to the origin. Then before we calculate the position of the indicator we translate its origin with this x and y. This gives us a map that always shows the user at its centre and the indicator at the correct position. To improve the map further we will add textures to the map, but before we can do that we need to generate the map in a way that keeps the rotation of the tiles consistent.

5.1.4. Generating the spanning tree

The map in Holonomy VR is generated using a spanning tree. This tree needs to uniquely identify each tile. Since we want to generate an order-5 square tiling, as seen in Figure 4.1, we use a few rules to accomplish this. From the origin, the tree starts in four initial directions: North, East, South and West. After the initial directions the tree continues to the left, forwards and to the right. However, these directions are not always allowed. The tree is constrained by the following two rules: The tree cannot take two consecutive right steps and the tree cannot take two left steps without a right step in between. These rules create a unique path for each tile. The generated tree path and an illegal path can be seen in Figure 5.3.

Using these rules we can generate the map. But we still need to know the hyperbolic coordinates of the generated tiles on the hyperbolic plane. In Holonomy VR the map was originally generated using the method described in a blog post from Martin Lu [33]. Lu uses an edge tree to calculate the positions of the corners of every tile. This is done by storing the coordinates of a corner in a polar transform model. The model can be described by three values. The first one describes the initial rotation, the second one describes the distance to the origin and the last describes the final rotation. Polar transforms allow for translating, rotating and transforming a polar transform by another polar transform. This method works excellently and is quite precise. However, the edge tree is not ideal since it forms a tree of edges and not tiles. So, to make the implementation of the map more flexible the edge tree needs to be related to the actual spanning tree path used in Holonomy VR. However, this would still generate the map from the origin from the tree path, without respecting the rotation of the tiles. This causes a problem for texture mapping. So this issue also needs to be resolved.

To solve these issues, the map is redesigned with the goal to be able to generate the map from any tile at its centre and use the spanning tree to keep the rotations consistent. This is done by generating

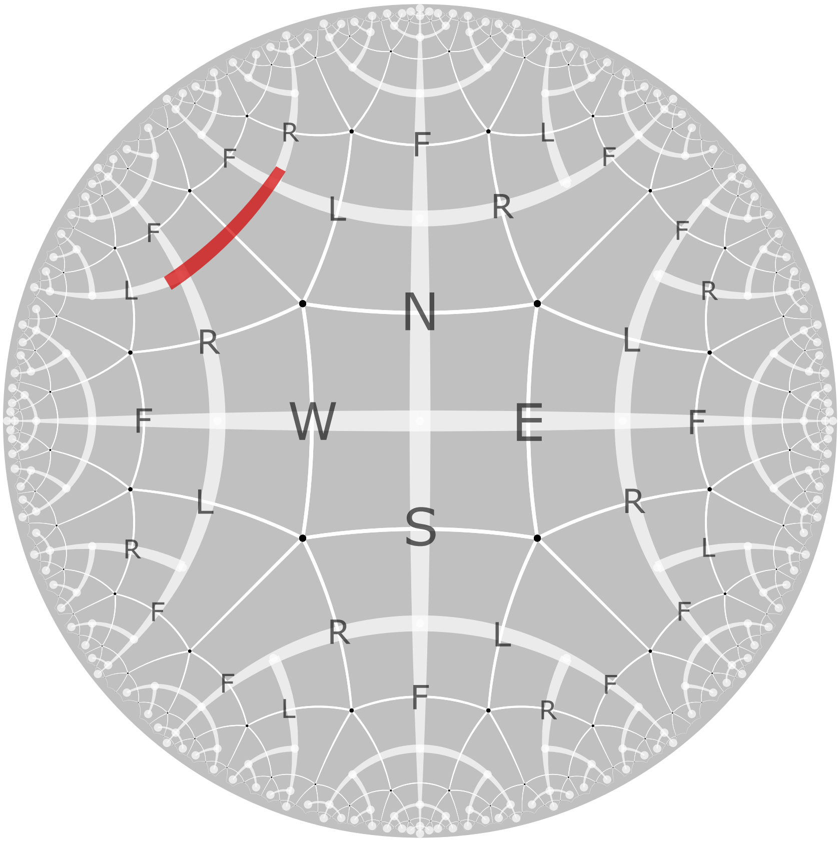

Figure 5.3: The spanning tree of the map with an illegal path in red. The initial directions are indicated with the letters N, E, S and W. The directions are indicated with the letters L, F and R.

the map once from the origin of the graph and, by not only storing the children of the tile, but also the parent of the tile. The generation directly calculates the location and rotation of the tile centres in hyperbolic space. But before these calculations can be done, the order-5 square tiling needs to be defined mathematically.

The order-5 square tiling falls under a regular hyperbolic tiling. These can be defined as (p, q)-tilings, where p is the number of sides of the polygon and q is the number of polygons that meet at a vertex. The order-5 square tiling is thus a $(4, 5)$-tiling. With this definition, the hyperbolic distance between the centres of the tiles can be calculated using the second law of the hyperbolic cosines [24]. This law is based on the fact that the lengths of sides of a hyperbolic triangle are determined by the interior angles of the triangle. This law has no Euclidean analogue since this would not be possible in Euclidean geometry. The law is defined as follows:

$$\cosh(\alpha) = -\cos(\beta)\cos(\gamma) + \sin(\beta)\sin(\gamma)\cosh(a)$$

Here the law is defined for a hyperbolic triangle with the sides a, b and c and the interior angles α, β and γ. To calculate the length of side a the law can be rewritten as:

$$\cosh(a) = \frac{\cos(\alpha) + \cos(\beta)\cos(\gamma)}{\sin(\beta)\sin(\gamma)}$$

For the $(4, 5)$-tiling the right-sided triangle as shown in Figure 5.4 is used to calculate the distances between the centre of two tiles. The general solution can be calculated with p and q as follows:

$$\alpha = \frac{\pi}{q}, \beta = \frac{\pi}{p}, \gamma = \frac{\pi}{2}$$

$$\cosh(a) = \frac{\cos(\frac{\pi}{q}) + \cos(\frac{\pi}{p})\cos(\frac{\pi}{2})}{\sin(\frac{\pi}{p})\sin(\frac{\pi}{q})}$$
$$= \frac{\cos(\frac{\pi}{q})}{\sin(\frac{\pi}{p})}$$

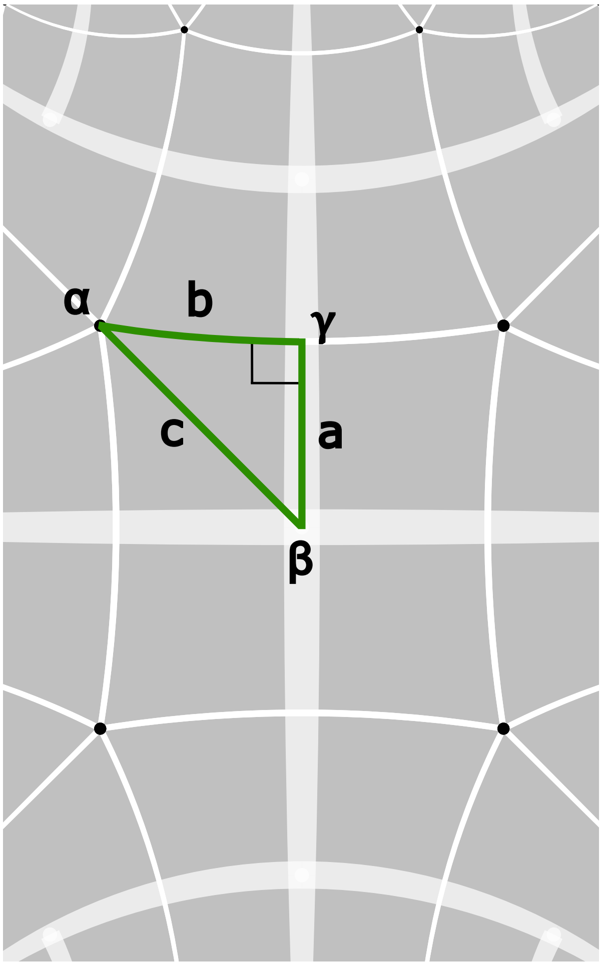

Figure 5.4: The triangle used to calculate the distance between the centres of the tiles. The distance between the centres is $2 * a$, the length of an edge is $2 * b$ and the distance between the centre and the corner is c.

The distances between the centres of the tiles are now known. The spanning tree and tile locations can now be generated. The tree is generated by starting at the origin and generating the first four tiles. The tree then continues by generating the children of the tiles up to a depth of 4. When the user moves to a different tile, we do not naively recompute the spanning tree and assign the correct tree paths, but we start generating the spanning tree from this new location, the user tile. First the user tile is shifted with a Möbius transformation to the origin, this keeps the current rotation of the user tile. Then we generate the spanning tree from the user tile following the original spanning tree. The tiles are generated while computing the difference in the update direction between the original parent and the current update direction. This sounds trivial but in practice is quite complicated, since we can also generate a tile from its child so the algorithm needs to be able to take a step backward.

There are a few different cases that influence this calculation. First of all it depends on if a newly generated tile is a centre tile, then the update is based on the initial directions and the spanning tree can generate the new children by rotating the initial directions and following them. Next the original tile can be the centre tile, in this case the tile is always updated from its parent and the update direction can be calculated from the initial directions. However, the generation needs to generate three children for this tile and its parent. These cases are trivial to find since the spanning tree of the order-5 square tiling has a maximum of one illegal direction. For other tilings this might not be the case and the algorithm needs to be altered to allow for this. But for this use case this algorithm will suffice.

The new tile in the spanning tree is now fully defined, its parent tile and its children are now all known. The only thing left for this new tile is calculating the centre transformation. This is done by first rotating the transformation of the tile that is used to update the tile in the direction of the new tile, then moving it by the centre distance and then rotating it to be in line with the original parent direction. Then the algorithm continues with the newly generated tiles until the end of the original spanning tree is reached.

Since we keep track of the parents of the tiles this algorithm does not require recomputing the tree paths to find its parents. So when the user moves to a new tile, the spanning tree only needs to calculate every tree path once. This is important since computing the tree path is one of the most computationally expensive parts of the generation.

Now that the centre locations of the tiles are generated, the actual corners of the tiles can be calculated. This can be done by using the same triangle as in Figure 5.4 and calculating c. Since a is known, c can be derived by using the law of sines [24]:

$$\frac{\sinh(a)}{\sin(\alpha)} = \frac{\sinh(b)}{\sin(\beta)} = \frac{\sinh(c)}{\sin(\gamma)}$$

$$\begin{aligned}\sinh(c) &= \sin(\gamma)\frac{\sinh(a)}{\sin(\alpha)} \\ &= \sin(\frac{\pi}{2})\frac{\sinh(a)}{\sin(\frac{\pi}{q})}\end{aligned}$$

Using c the corners are generated based on the centre of the tiles. This gives us the basis for the map, the basis is only updated when the user moves to a new tile. After every update the representation is stored and sent to the TextureData class which handles the movement of the player and the rendering of the map. Since we have now created a tile tree that keeps the rotation of the tiles consistent we can discuss the next improvement, texture mapping.

5.1.5. Texture Mapping

In Euclidean geometry, textures can be mapped to triangles using barycentric coordinates. We can try the same method in the Poincaré disk model to calculate these coordinates, however, this method does not take the curvature of the space into account. Nonetheless, with some changes, we can use this method also in hyperbolic geometry. By using hyperbolic barycentric coordinates [3] we can map textures to hyperbolic triangles.

The hyperbolic barycentric coordinates can be found by projecting the hyperboloid model onto the tangent plane. This plane gives us a relation to the planar barycentric coordinates. The planar and barycentric coordinates are related to each other. Both of them can be calculated using the area of the triangles, however, for the hyperbolic barycentric coordinates we also need the length of the edges.

This means that we could still use the Poincaré disk model, however, calculating the area of a triangle in this model involves calculating the inner angles of the triangle which requires shifting every point of the triangle to the origin which is not efficient.

By dividing the hyperbolic squares into hyperbolic triangles we can calculate the planar barycentric coordinates from the hyperbolic model to map a Euclidean texture to the hyperbolic square. For every type of tile, we have made a texture to represent it in 2D. By applying this to the map we add more detail which allows the user to relate the map to the virtual environment. This detail is also necessary for the 2D version of Holonomy VR which will be discussed in subsection 5.1.7. An example can be seen in Figure 5.5.

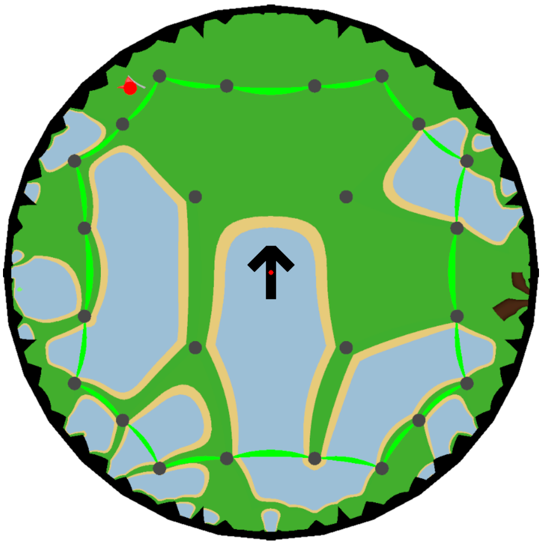

Figure 5.5: The Holonomy VR virtual environment transformed to a 2D experience

5.1.6. Indicator

The indicator is a feature of the map that is implemented as a red dot that indicates where the objectives are located. This implementation was also based on the edge tree mentioned in 5.1.4 which meant that the indicator needed to be reimplemented. When the objective is outside of the map the indicator will render at the edge of the map. When this is the case it will still use the original method to calculate the path from the user to the objective however the transformations are now based on the tree path itself using the polar transformations. The resulting polar transformation gives the angle from the user to the objective. This is then indicated by placing a point on the edge of the map in that direction.

When the objective is inside the map the indicator will render at the correct location. This is done by placing the indicator on the centre of the tile which contains the objective. To keep the indicator centred when the user moves, the indicator is transformed with the same Möbius transformations as discussed in subsection 5.1.3. This gives the correct position of the indicator.

The user can also create their own indicators to keep track of visited locations. This indicator is created at the user location and is displayed as a blue dot on the map. The rendering of the user-created indicator works the same as the objective indicator. The user is able to remove the indicator by moving to the location of the indicator.

5.1.7. 2D View

Now that the map is accurate, detailed and always related to the player position we have transformed it into a 2D experience. This lets users explore Holonomy VR in 2D. This is done by rendering only the map in full screen. The map is generated in the same way as in VR and the user can move around by using the arrow keys. To ensure that users have the same movement limitations, borders are used that represent the physical borders of the original VR environment. This view still allows users to see tiles beyond the borders, but we can also disable them to give them the same information as in VR.

The 2D view can also be used to mimic looking at a map for navigation purposes. This is done by allowing the user to collect objectives by being in a location close to the objective. The objective is then marked as visited and the user can move to the next objective. This means that users can then complete objectives that are normally beyond the border. The goal of this mode is that users can learn the layout of the environment and use this knowledge to navigate in VR without having to actually reach all the objectives in the training stage.

5.2. Other Additions

Besides the map, Holonomy VR has gotten some other new features that are important to mention. These include beacons, landmarks, sounds, result tracking and the addition of a user study menu.

These features are added to improve the user experience and to allow for user studies to be conducted.

5.2.1. Beacons

Beacons are an extension of the indicators used by the map. They are displayed in VR behind the bushes and are used to guide the user to the objectives. The basis of its position is calculated in the same way as the indicator position. It is done by calculating the angle between the user and the target, then by calculating the hyperbolic distance we have the correct hyperbolic position of the beacon. However, unlike the indicator, the beacon must be placed in VR which means that the position must be transformed to a logical position in the VE. This is done by scaling the distance of the beacon such that the user can relate the distance to the beacon with the rough distance to the objective. This is not a perfect solution, but it indicates the general direction and distance to the objective. An example of the indicator and beacon can be seen in Figure 5.6.

An indicator placed by the user will be visible as a beacon, when the indicator is inside of the reachable area the beacon is displayed inside of the bushes. The beacon will be displayed in the correct location by placing the beacon on the correct tile and setting rendering layers. To keep the indicators and beacons consistent they are shown in the same colours.

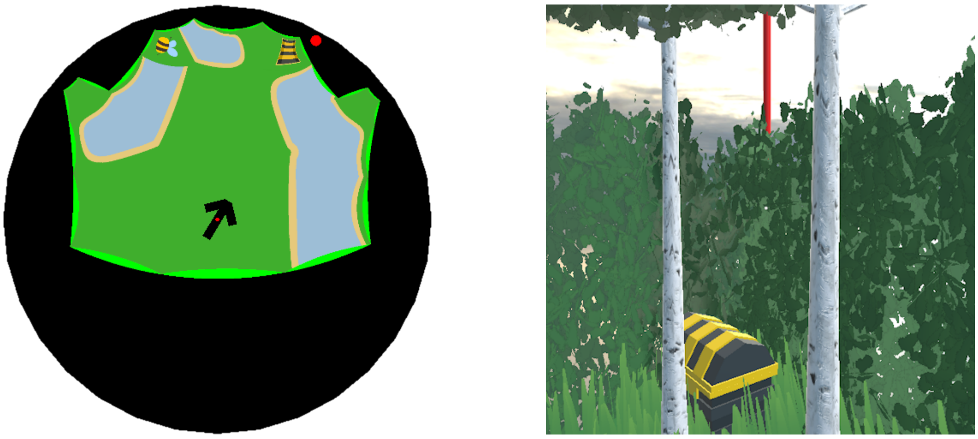

(a) 2D indicator **(b)** VR beacon

Figure 5.6: Guidance to an objective that is outside of the grid

Another addition that comes with beacons are sounds. When an objective becomes directly reachable a positive sound will play. When the objective is no longer reachable a negative sound will play. This is done to give the user feedback on their actions. A sound will also play when the user reaches an objective. A priority system is used to make sure that more important sounds overrule other sounds. For example, an appearing sound is more important than a collection sound because the user otherwise might not know that another objective is now reachable.

5.2.2. Landmarks

Originally Holonomy VR had two different game modes, finding the flag and collecting keys to open a chest. However, these game modes are not sufficient for research into navigation. So, a new game mode is added where the goal is to find landmarks in a specific or non-specific order. This game mode is used as a navigational task, since landmarks are used by humans to create cognitive maps and allow them to navigate more efficiently in an environment.

The available landmarks can be viewed in Figure 5.7 and are placed in Holonomy VR in the same way as the objectives. All of the landmarks have been designed to be recognisable in VR [49, 26, 25, 13, 2, 1, 5] and in 2D. The user needs to find all the landmarks in a given or in a random order. They can be collected by being in the same tile as the landmark. The user can also see the landmarks on the map and the 2D view. When the game mode is executed in the unordered mode users will be able to find their way back to already found landmarks. The beacons and indicators for these landmarks will turn green. This allows them to easily revisit landmarks if they want to.

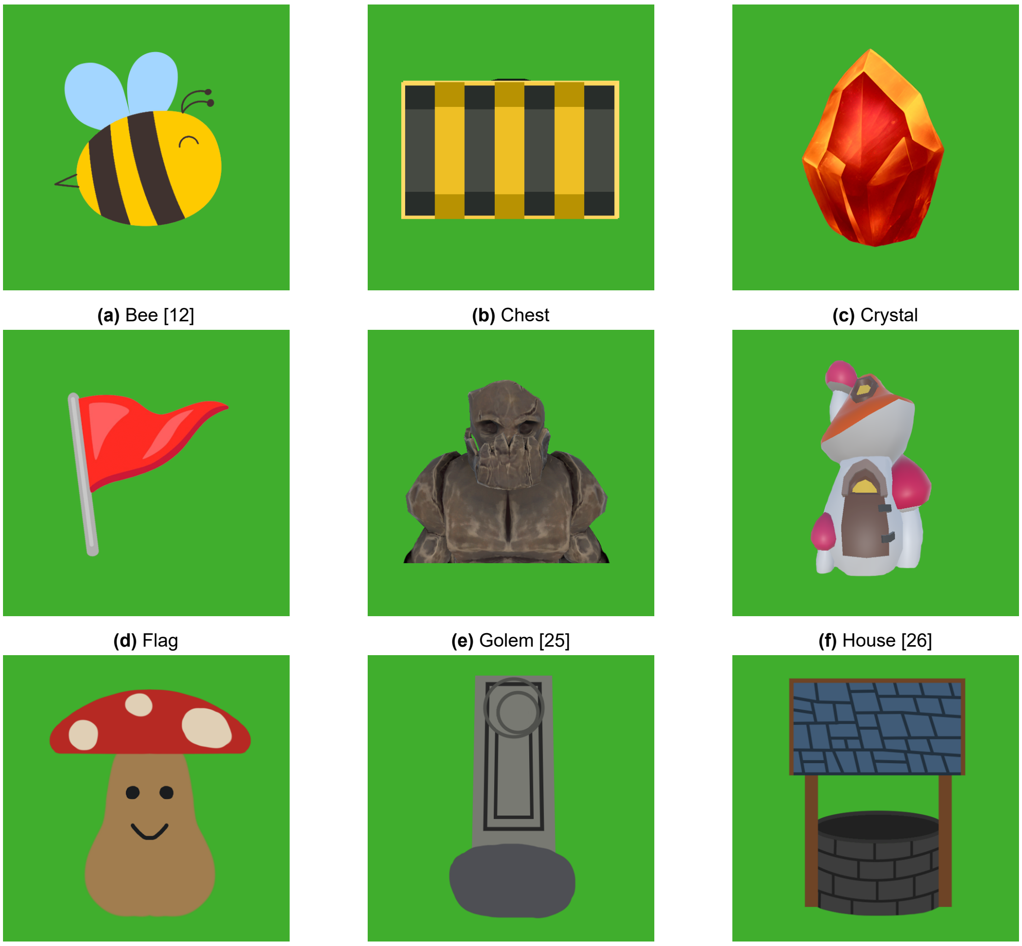

(a) Bee [12] **(b)** Chest **(c)** Crystal **(d)** Flag **(e)** Golem [25] **(f)** House [26] **(g)** Mushroom **(h)** Stone **(i)** Well

Figure 5.7: The available landmarks in Holonomy VR

5.2.3. User Study Capabilities

To help us execute a user study the user study menu is created. This menu allows the user to select their participant ID, group and the stage they need to complete. Their information is automatically saved so they can continue quickly to the next stage. While they are doing a stage their progress is being tracked and automatically saved when the stage is completed. This data is used to analyse the performance of the user. This data includes the information of the user, which objectives and in which order the objectives were visited, the time the user took and the path the user took. These results are saved in a CSV file on the VR headset or on the computer when using the 2D view.

The groups in the menu are defined in code using different settings. These settings are used to create different conditions for the user study. As a general improvement these settings can be changed in the settings menu. This allows the user to create an experience tailored to their needs. Both of the menus can be seen in Figure 5.8. In the future it would be good to be able to load the group settings and levels used during a stage from a file such that any researcher can create their own user study.

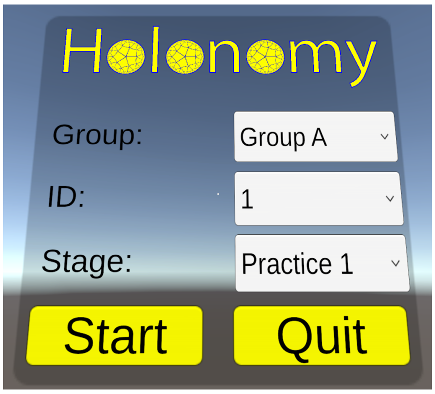

(a) The user study menu

(b) The settings menu

Figure 5.8: The user study and settings menus

6
Method

6.1. Setup

The setup of the study is designed around the hypotheses presented in section 4.2. The study is divided into three parts: gathering information about the participant by means of a survey, doing tasks in Holonomy VR, and continuing the survey by answering questions about the locations of landmarks encountered in the tasks and questions about hyperbolic geometry. The complete survey can be found in Appendix B, where applicable the correct answers are stated. Before the study starts, the participant is asked to sign an informed consent form. This form is used to inform the participant about the study and to get permission to publish the data gathered during the study. The informed consent form can be found in Appendix D.

6.1.1. Survey Part One

The first part of the survey is designed to gather information about the participant. The survey consists of questions about the age, gender, field of study[1], educational attainment, and dominant hand. This information is necessary to relate the performance of participants to their characteristics as discussed in subsection 3.1.5. Next, the participant is asked about their experience with video games and virtual reality. These experiences might influence the performance of the participant in the tasks even though the tasks are designed to be as simple and understandable as possible. Then we ask the participant about their current understanding of hyperbolic geometry. This is used to compare the knowledge of hyperbolic geometry before and after the tasks. These last three questions are answered using a 6-point scale.

The next two sections of this part of the survey are about navigation and exploration. These are related to the creation of the cognitive map and to analyse the differences in exploration tendencies between the participants. To analyse the navigational abilities, the Santa Barbara Sense-Of-Direction Scale (SBSOD) [22] is used. To analyse the exploration tendencies, the Curiosity and Exploration Inventory-II (CEI-II) [27] is used. The SBSOD measures the sense of direction. The CEI-II is used to measure the exploration tendencies of the participant. The results of these two sections are used to see if the navigational abilities and exploration tendencies of the participants influence the performance in the tasks. This is mostly important for the psychological research that Leiden University is conducting. After completing the first part of the survey, the participant will continue with the tasks in Holonomy VR.

6.1.2. Hyperbolic Experience

To start, the participants will be divided into four groups, each group will receive training using different methods, except for the control group. The groups are as follows: No training (Control), training in virtual reality (VR), training in virtual reality with a map (VR-Map), and training in 2D with only the map (Map-Only). VR-Map has access to the map on their wrist, which is shown in Figure 6.1. However, this map does not show the bushes and trees, this was done to be able to see the differences in the

[1] Gathered using the ISCED 2013 fields [55]

affordance that this information on the map gives.

The groups will do three practice tasks, the training task, except for the control group, and two evaluation tasks. The layout of the tasks is the same for each participant to ensure that the results are comparable. To ensure that the same explanation and setup are used for each participant, a script is used. The script can be found in Appendix C.

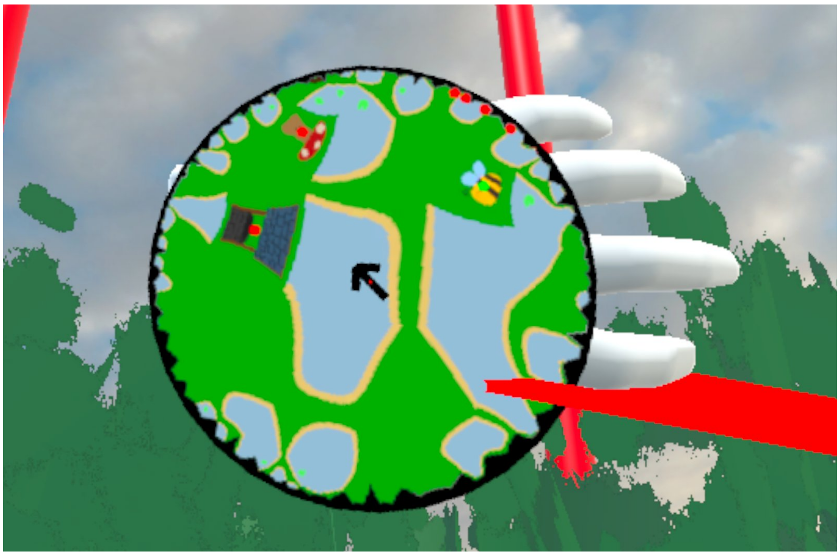

Figure 6.1: The map on the wrist of the participant

The goal of each task is to find one or more landmarks. During the task, each participant will have access to the beacon when they are in virtual reality, when using the 2D experience they will have access to the indicator which corresponds to the location of the beacon as discussed in subsection 5.2.1. Depending on the group, the participant will have access to the map during the practice and training sessions.

The practice stage is designed to get the participants familiar with Holonomy and the cues that are presented to them. The training stage is designed to train the navigational ability of the participants by giving them the task to find eight landmarks and remember their locations. The evaluation stage is designed to test the navigational knowledge of the participants to see if they can now navigate more efficiently to the landmarks. The goal of each evaluation task is to find four landmarks in a given order. During the stages, the performance of the participants is recorded by Holonomy VR. All stages have a set time limit for each session which makes sure that the study will not take too long when the participant is not able to complete the task. The stages will be discussed in more detail in the following sections.

Practice stage

To make sure that all participants get familiar with Holonomy VR and the accessible cues, participants will be trained using three practice sessions with the goal to find the flag. For the Map-Only group, the practice sessions are also used to get familiar with the controls. In each session, the flag is moved further which increases the difficulty per task. The first session is used as a way to give an explanation to the participants, this allows them to directly relate the explanation to what they see. First, they are told what the trees (or dots in 2D) and bushes (or walls in 2D) are, what the beacon (or indicator in 2D) represents, how the beacon and its sounds work and if applicable how they can access the map and how they can move around. It might seem to the participant that they are in a closed-off space because they are surrounded by bushes. However, they are told that this is not the case since they are in a hyperbolic world and that walking around a tree will rotate the world which will move the bushes and allow them to explore different parts of the world. After the explanation, the participant can easily walk to the flag since the flag is placed one step to the north, and one step to the left. This means that the flag is already in sight and directly reachable by the participants.

In practice session two, the flag is placed one step north and a step forward. This places the flag just outside of the direct reach of the participant. This means that the participant has to walk multiple times around a tree to reach the flag. Here the participant has to learn on their own how to explore previously unreachable locations. This allows the participant to adapt to the hyperbolic world and learn how to navigate in it. This relates to the first hypothesis. The third practice task is similar to the second practice task, however the flag is placed an extra step forward. This allows the participant to get more practice in navigating in the hyperbolic world. By now the participant should be able to navigate and be familiar with the cues that are presented.

Training stage

In the training task, the participant is asked to find eight landmarks and to remember their location. The order in which the landmarks need to be found does not matter. The landmarks are placed in different locations from the starting point, an overview of the landmarks is shown in Figure 6.2. The landmarks are spread in all directions and placed at most four tiles from the starting point. This requires the participant to explore the environment to find all landmarks. The Mushroom and the Well are only separated by one tile, this makes sure that the participant sees both of these landmarks at the same time. This is used as a control question when we ask the participant about the locations of the landmarks.

The Map-Only group has the advantage that they only need to get close to the landmark to collect it. This is done to mimic the experience of looking at a map. This means that this group can also collect a landmark when it is just outside of the border.

By recording the time taken to find the landmarks we can see if the training session has an effect on the navigational ability of the participants. This relates to the second hypothesis. After completing the training, the participant should be able to recollect the locations of the landmarks and navigate to them more efficiently.

The control group does not take part in the training session. Instead, they will do a distraction task. This task is designed to keep the participant occupied to make sure that they spend the same amount of time on the training session as the other groups. The task consists of playing the computer game pong for a set amount of time, or by winning two times. By not giving the control group a training session we can see if the training session has an effect on the navigational ability of the participants. This relates to the third hypothesis. After completing the distraction task, the participant continues with the evaluation sessions.

Evaluation stage

The evaluation sessions are designed to test the navigational ability of the participants. The goal of each evaluation task is to find four landmarks in a given order. The landmarks are placed in the same location as in the training, also all landmarks from the training are present. Now the order in which the participant needs to find the landmarks is given. This allows us to compare the navigational performance of the participants. The time and path taken to find the landmarks are recorded to see if the training task has an effect on the navigational ability of the participants. This also relates to the third hypothesis. Using this data we should also be able to verify the fourth hypothesis.

Since the Map-Only group will now experience Holonomy VR for the first time in virtual reality, this group will get instructions on how the 2D experience translates to the VR experience. The participants are told that the grey dots correspond to the trees and that the green boundary corresponds to the bushes. Since the participant can move by walking it is expected that they do not need to practice in VR beforehand. After completing the evaluation tasks, all participants continue with the survey which will be discussed in the next section.

6.1.3. Survey Part Two

The last part of the survey consists of three sections. The first section is about the locations of the landmarks that the participant encountered in the tasks. The participant is asked four times to mark the landmark that is the furthest away from the three given options. This is done to see if the participant has created a cognitive map of the hyperbolic world. The second section is about the knowledge of hyperbolic geometry. The participant is asked a series of questions about the properties of the hyperbolic geometry that they experienced. These questions are designed to see if the participant has picked up on the properties of the hyperbolic geometry. This indicates if the affordances presented

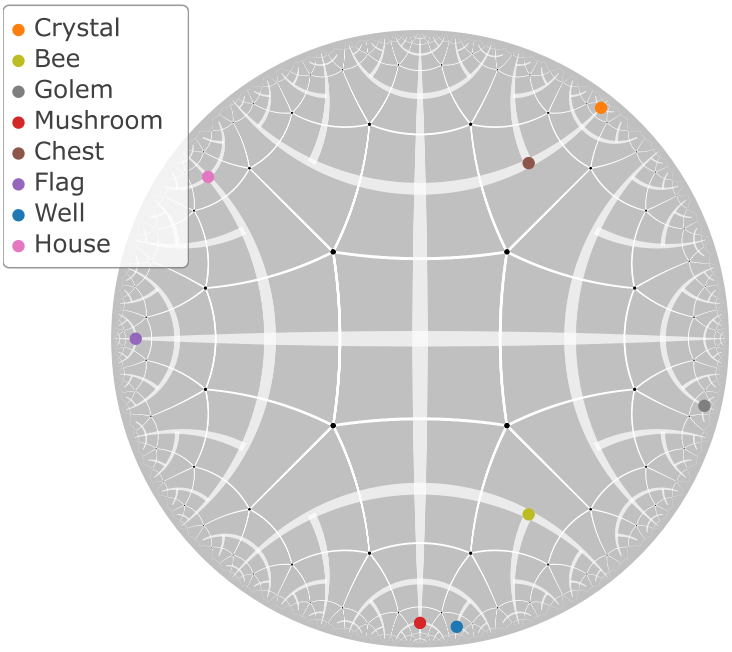

Figure 6.2: The locations of the landmarks

during the tasks translate to the properties of hyperbolic geometry. This allows us to verify the fifth hypothesis. Since the user should now understand the hyperbolic properties we can relate that back to the navigational performance, verifying the sixth hypothesis. The last section consists of general questions: "What did you think of the experience?", "How do you evaluate your performance for the evaluation task?", "Did you learn anything about hyperbolic geometry?", "Do you have anything else to add?". These questions are used to validate the found results with the thoughts of the participants. After completing the survey, the participant has finished the study.

6.2. Pilot User Study

To verify if the setup of the study is plausible a pilot user study is conducted. It was assumed that a study session would take no more than 90 minutes. The pilot user study is done with a group of 9 students aged between 18 and 32. One of them is a woman, the others are men. The pilot was held at the office of First Element[2]. The participants were asked to fill in the survey and do the tasks in Holonomy VR. The participants were equally divided over the four groups, except the Map-Only group had one participant more. In the room, we measured an area of 3 by 3 meters to which we set the boundaries of Holonomy VR. An overview of the room is shown in Figure 6.3a.

In the pilot, it was found that most participants finished within 45 minutes and that the time slots could thus be reduced to 60 minutes. 66% of the participants managed to finish all or nearly all stages, which is an acceptable percentage. However, it was found that the two evaluation sessions were not equal in length. It would be better to make the evaluation sessions equal in length to make sure that the results are comparable. A workaround would be to compare the results by using the optimal path for the sessions. Then we could compare them using a ratio. Overall, the pilot was a success and the setup of the study was found to be successful.

6.3. User Study

There were a few improvements we could make to the study from what we have learned from the pilot. However, since many participants are needed to get potentially significant results, we opted to continue

[2] `First Element` is a company that focuses on providing geographic insights to businesses, they were kind enough to lend out a room in their office for the study.

(a) Office at First Element **(b)** Sylvius VR Lab at Leiden University

Figure 6.3: The user study locations

with the same setup as the pilot. This means that we can include the data of the participants from the pilot in our results. In the overview of the participants that follows, the data of the pilot participants are also included.

The study was conducted in Delft at the same location as with the pilot and at Leiden University. The setup used in Delft and Leiden can be found in Figure 6.3. Both locations used an area of 3 by 3 meters to set the boundaries of Holonomy VR. To make sure that the participants were given the same information by both researchers, a script was used as mentioned previously. The script can be found in Appendix C.

In total, 66 participants took part in the study. However, 4 of them were measured incorrectly and are thus excluded from the results. For another participant, the screen casting from the headset was not working, so no hints could be given. However, the participant performed well so we assume that hints were not necessary for this participant and thus this participant is still included. This results in a total of 62 participants.

The participants were equally divided over the four groups. The control and VR-Map group had 16 participants and VR and Map-Only group had 15 participants. The participants were aged between 18 and 34 and all studied or did study at the TU Delft or Leiden University. 24 of them are women, 37 are men and one is non-binary. The control and Map-Only group contain more men than women, the others were well balanced. 27 participants finished their high school education, 27 participants have a bachelor's degree and 8 have their master's degree. Most of the participants studied "Engineering, manufacturing and construction" or "Social sciences, journalism and information". The participants were asked to fill in the survey and do the tasks in Holonomy VR. The results of the study will be presented in the next chapter.

7 Results

In this chapter, we present the results from the user study that was conducted. We will first take a look at the completion ratio of the participants. Then we dive deeper into the performance of the completed stages and note interesting findings between the groups. Next, we present the results from the questions about the landmarks and the hyperbolic properties and finally, we list a summary of the open questions. The values from the graphs presented can be found in the appendix in Appendix A and each group is represented by a different colour which is also colour-blind safe[40].

7.1. Completion

We first talk about the completion of the stages. A stage is completed when the participant has collected all landmarks within the given time. From Table 7.1 we can see that about 65% of all stages were completed. This does not differ significantly between the groups, except for the training stage. Here Map-Only completed 100% of the training, while the other groups are less successful.

Group	P 2 Mean	P 2 SD	P 3 Mean	P 3 SD	T Mean	T SD	E 1 Mean	E 1 SD	E 2 Mean	E 2 SD	Total Mean	Total SD
Control	62.50%	50.00%	81.25%	40.31%	-	-	47.06%	51.45%	35.29%	49.26%	57.81%	33.81%
VR	80.00%	41.40%	73.33%	45.77%	40.00%	50.71%	66.67%	48.80%	46.67%	51.64%	61.33%	32.48%
VR-Map	68.75%	47.87%	81.25%	40.31%	50.00%	51.64%	64.70%	49.26%	58.82%	50.73%	65.88%	37.26%
Map-Only	86.67%	35.19%	86.67%	35.19%	100.00%	0.00%	57.14%	51.36%	47.14%	51.35%	76.00%	17.24%

Table 7.1: Ratio of completion by stage. ”P” stands for practice, ”T” for training and ”E” for evaluation.

Next, we dive deeper into the completion of the stages. The completion can also be calculated by taking into account the individual landmarks that need to be collected and grouping it by the sessions. For the practice session, we only take into account practice stage two and three, since an explanation was given in practice stage 2 and the participant could directly walk to the target. In the practice session, participants needed to collect a total of 2 landmarks, in the training session and evaluation session they needed to find 8 landmarks in both. The percentage is calculated by dividing the number of landmarks collected with the total that could be collected. This is shown in Figure 7.1. Here it is shown that participants performed better than by purely looking at the completed result alone. However, the groups still do not differ significantly from each other, except again for the training with only the map.

7.2. Performance

Since the completion rate is not 100% for each group we will only take the completed runs into account. This allows us to display the results in more detail. First, we take a look at the time spent in each stage. These results are shown in Figure 7.2. Again in most stages, there is no significant difference between the groups, however, when we take a look at the training stage we can clearly see that Map-Only performs much better than the others.

Another performance metric that is interesting is the total steps taken per stage. This is shown in Figure 7.3. Here we can see that Map-Only in practice 3 has some large outliers. Another high outlier

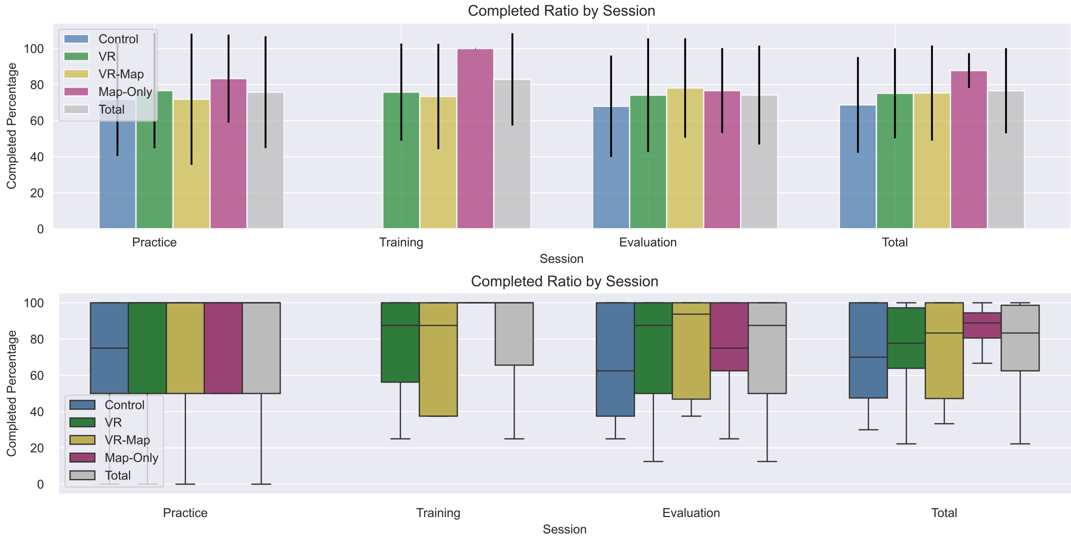

Figure 7.1: Percentage completed per session

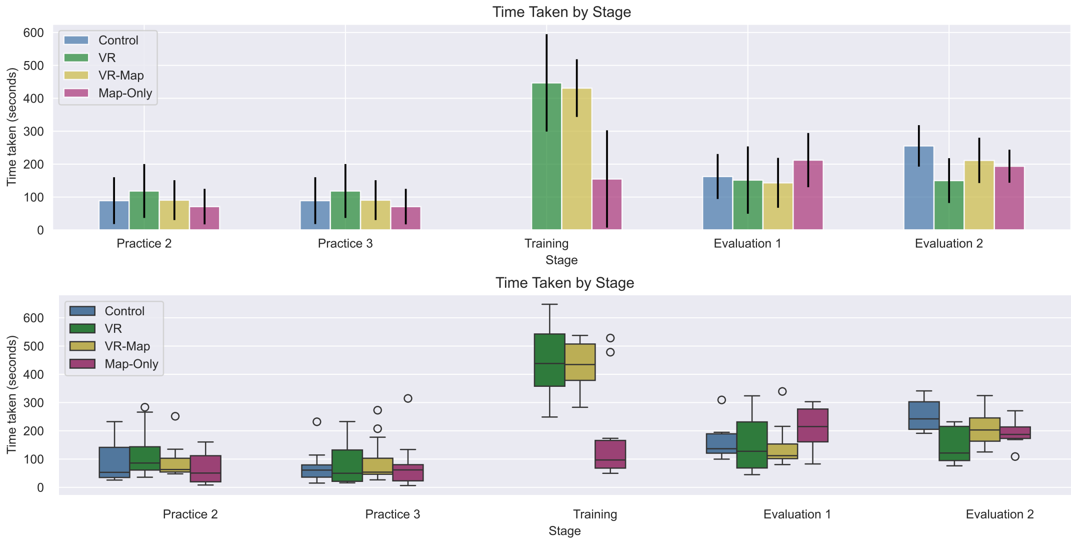

Figure 7.2: Time taken for each stage

is in the training stage, but Map-Only still shows a better performance. However, it would be good to take a closer look at this outlier. The outlier is participant 36, who got lost during the training stage. This participant reached 14 steps away from the origin, while the mean for this group is 5.67 and the standard deviation is 3.06. The path taken by this participant can be found in Appendix A, Figure A.1. For the progression results, this participant will be excluded.

7.3. Progression

Next, we take a look at the progression of the participants within the stages. First, let's take a look at the time when the participant collected each landmark. This is shown in Figure 7.4. We start with the training. Here it can clearly be seen that Map-Only collects the landmark faster and with fewer steps. This is not the case for the evaluation session. Here we see that Map-Only takes longer to collect the first landmark in the first evaluation stage. Another interesting thing to note is that the control group seems to take a longer time to collect the first landmark in the second evaluation stage. When we combine the training and evaluation sessions we see that Map-Only starts off faster but then slows down when they get to the evaluation stages. However, they are overall still faster. We can also take a look at the standard deviation. These are larger in the evaluation stages, especially in the second

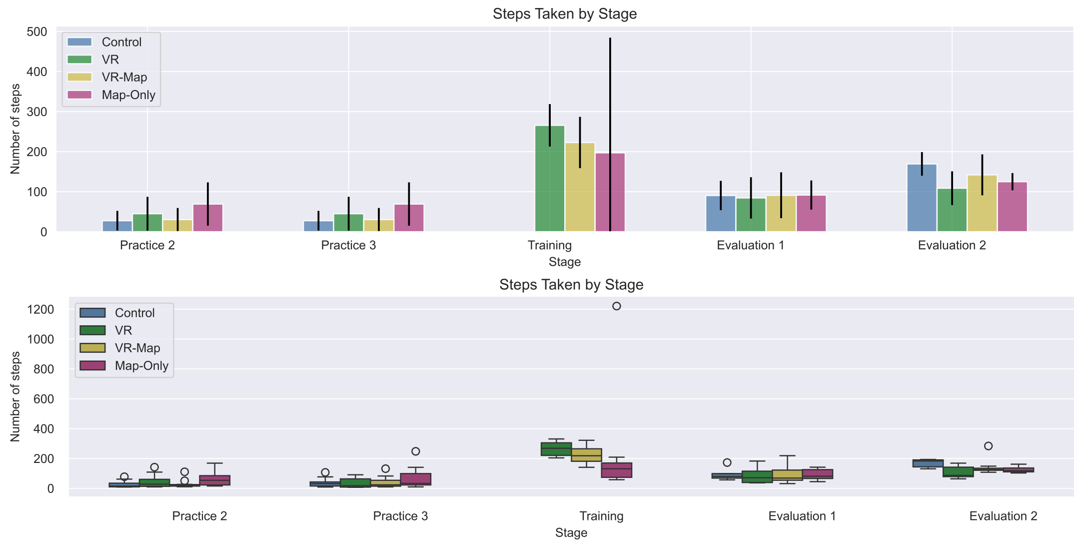

Figure 7.3: Total steps taken for each stage

evaluation stage.

(a) Collection time and steps during training

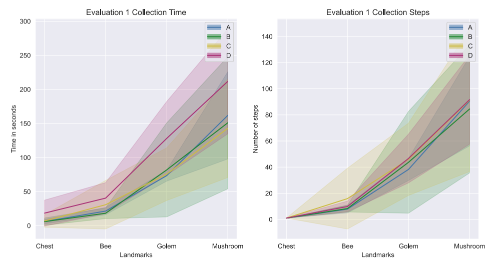

(b) Collection time and steps during Evaluation 1

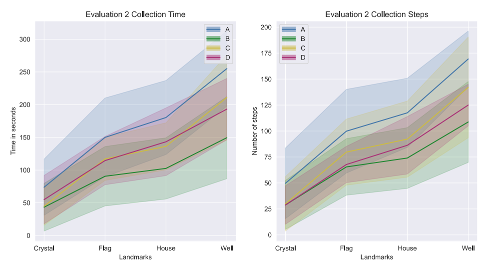

(c) Collection time and steps during Evaluation 2

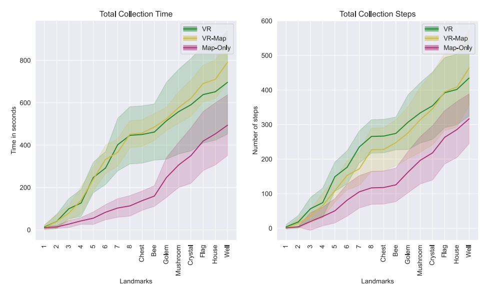

(d) Total Collection time and steps

Figure 7.4: Collection time and steps taken for each landmark

Figure 7.5 shows the percentage completed of the stage in relation to the total distance travelled by the participant. The plots can display differences in reaching landmarks by comparing the number of steps. In the training stage, it can be seen that Map-Only collects the first landmarks faster than the other groups. In the evaluation sessions, all groups perform the same, but then when we look at the total progression we see that Map-Only is faster at first but then takes longer. This makes sense since, in the training, the Map-Only group only needed to be near the landmark to collect it. So we indeed see that this is faster than needing to reach the actual landmark.

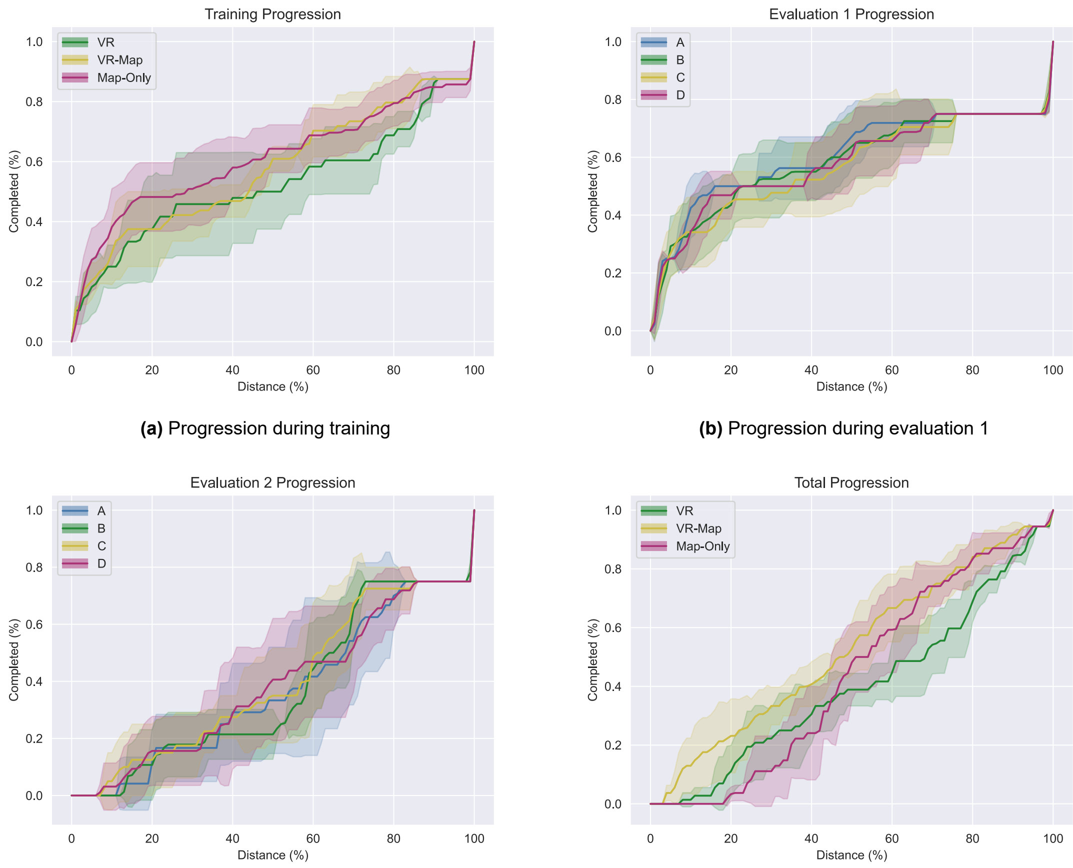

(a) Progression during training

(b) Progression during evaluation 1

(c) Progression during evaluation 2

(d) Total progression

Figure 7.5: Progression data

7.4. Answers to Landmark and Hyperbolic Questions

First the landmarks, these results are displayed in Table 7.2. It can be seen that participants found it difficult to answer the second and fourth questions about the landmarks. Only the control group answered the first question relatively well, the third question was answered correctly by all of the groups, VR and Map-Only answered 100% correctly. But overall the results indicate that participants found it difficult to answer the questions.

Group	1 Mean	1 SD	2 Mean	2 SD	3 Mean	3 SD	4 Mean	4 SD	Total Mean	Total SD
Control	37.50%	50.00%	6.25%	25.00%	81.25%	40.31%	31.25%	47.87%	39.06%	18.18%
VR	60.00%	51.71%	26.67%	45.77%	100.00%	0.00%	46.67%	51.64%	58.33%	20.41%
VR-Map	31.25%	47.87%	25.00%	44.72%	75.00%	44.72%	18.75%	40.31%	37.50%	15.81%
Map-Only	46.67%	51.64%	6.67%	25.82%	100.00%	0.00%	53.33%	51.64%	51.67%	19.97%

Table 7.2: Percentage correctly answered for each question about the relative distance between landmarks

Next, the results from the questions about the hyperbolic properties. These can be found in Table 7.3. Here we see that participants found it difficult to answer questions one and four. However, participants performed very well on question two. Control and Map-Only did also quite well on question three. But overall the results indicate that not every participant had learned about the hyperbolic properties.

7.5. Open Questions

At last, we present the open questions. First of all, what did the participants think of the experience? Most of the participants found the experience fun and interesting. Although they also found it challeng-

Group	1 Mean	1 SD	2 Mean	2 SD	3 Mean	3 SD	4 Mean	4 SD	Total Mean	Total SD
Control	25.00%	44.72%	75.00%	44.72%	62.50%	50.00%	56.25%	51.23%	54.69%	29.18%
VR	26.67%	45.77%	93.33%	25.82%	40.00%	50.71%	33.33%	48.80%	48.33%	27.49%
VR-Map	12.50%	34.16%	87.50%	34.16%	43.75%	51.23%	50.00%	51.64%	48.44%	21.35%
Map-Only	26.67%	45.77%	100.00%	0.00%	86.67%	35.19%	26.67%	45.77%	60.00%	15.81%

Table 7.3: Percentage correctly answered for each question about the hyperbolic properties

ing and sometimes confusing. Some participants who were not successful stated that they found it frustrating since they did not grasp the concept.

Next, participants seem to have a good understanding of their own performance. People who performed well also thought that they performed well and people who did not, said that their performance was lacking.

We also asked the participants if they learned something about hyperbolic geometry. For the people who could navigate effectively, they still found hyperbolic geometry confusing. But some also stated that they found it interesting and that they would like to learn more about it. Some people also stated that they found it difficult to navigate in hyperbolic space but were becoming better at it. When participants could not navigate effectively, they stated that they did not learn anything.

8
Discussion

In this chapter, we will discuss the hypotheses and relate them to the results of our study. To draw conclusions, we will look at the significance between different groups and the correlation between different variables. The significance will be tested by using the t-test, and we call it significant when it results in a p-value lower than 0.05. This, together with a high or moderate correlation, allows us to draw conclusions about the hypotheses. For the multiple-choice questions, we will compare them to the chance that they can be guessed by looking at the binomial probability. We will also look at results that are not directly related to the hypotheses. The hypotheses are divided into their relation with navigation and the hyperbolic properties. We will start with the navigation hypotheses.

8.1. Navigation

We start with the first hypothesis. We say that a participant can navigate when 75% or more of the evaluation is completed; we then say that it "clicked", or that they understood how to navigate. The completion percentage can be viewed in Figure 7.1. This method gives us a chance to include participants who might have gotten lost at some point, which makes completing the stage more difficult, and because of that, ran out of time. We could also use the notes about the participants' performance that were taken. However, to make it consistent, we opted for a more scientific method. This method corresponds well with the notes that were taken.

It clicked with about 60% of the participants in all groups, so they were able to navigate on the hyperbolic plane. Unfortunately, this result indicates that not everyone could successfully navigate. In Figure 7.5, we do see that all groups progress consistently, except for Map-Only in the training session. This group progresses relatively slower since they can reach many landmarks quite fast, while the others take more time to reach.

The relatively low completion rate might have multiple issues. First, the practice sessions might have been too easy; however, the completion ratio between the sessions is not significantly different. The data does indicate that the type of study might have an influence on the results. Participants that study "Social sciences, journalism and information" do seem to perform worse than participants who studied "Engineering, manufacturing and construction", "Information and Communication Technologies" and "Natural sciences, mathematics and statistics". The participants from these studies perform significantly better. This difference is mostly caused by the moderately negative correlation of -0.35 caused by the "Social sciences, journalism and information". This might be because the positively correlated studies provide more technical knowledge. Another cause could be that participants got lost; this can be explained by the maximum distance reached from the origin. The difference between the participants that have completed the stages and participants who did not is significant with a p-value of 0.000 and a correlation of 0.49. The difference in the number of steps can clearly be seen in Figure 8.1.

Even though hints were given if the participant wandered away from the target, it seems that most participants were not able to find their way back when it happened. This also correlates with the researcher that supervised the participant; however, this could be caused due to the different studies, since the research at Leiden University mostly had students from the "Social sciences, journalism and information". However, it would have been good to have automated this hint to make sure the hint was always given under the same conditions. Nonetheless, we can conclude that at least some people can adapt to navigate in hyperbolic space without any specific explanation about the hyperbolic properties. This allows us to move to the next hypothesis.

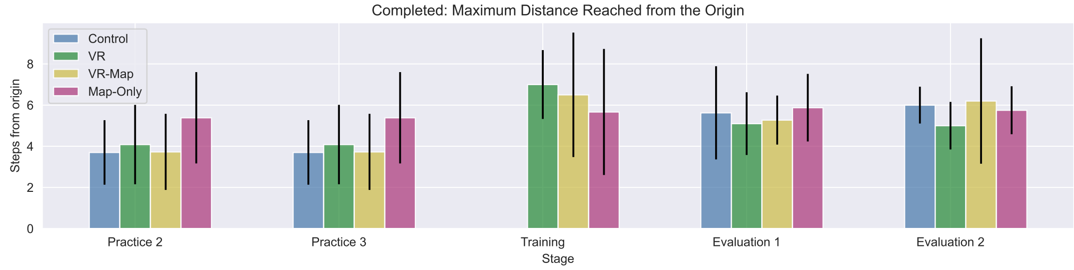

(a) Stages that have been completed

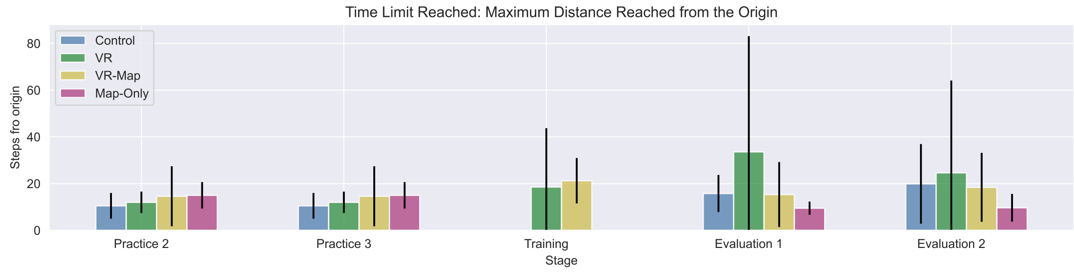

(b) Stages where the time limit is reached

Figure 8.1: Maximum number of steps reached from the origin

The second hypothesis is about the time taken to complete the training session. Here we only look at the completed training sessions. From Figure 7.2 and from Figure 7.3 we can clearly see that Map-Only performs faster. Especially when removing the outlier, which is explained in chapter 7. We can also see a clear difference in the time taken to collect each landmark when looking at Figure 7.4a. The difference between Map-Only and the others is significant. The p-value is less than 0.001 when comparing Map-Only with the other groups. Map-Only also has a correlation of -0.80 with the training time and a correlation of -0.74 with the number of steps taken.

This is what we expected since the training session of Map-Only is easier because they only need to get close to the landmarks. However, there is no significant difference between VR and VR-Map. This is not what we expected since we hypothesised that it would be easier when the participant has access to the map. This might be because VR-Map takes more time to look at the map and think about their movements. They seem to take fewer steps, but this is not a significant difference. It could also be caused by the lack of borders and trees displayed on the wrist map. So we can only conclude that the 2D Experience results in a faster training time than the VR experience with or without access to the map.

Next, we look at the third hypothesis. This hypothesis is about the time taken to complete the evaluation when a form of training was taken. Again when looking at Figure 7.2 and Figure 7.3 we see that there is no significant difference between the groups that have taken the training. This is also confirmed by Figure 7.4b and Figure 7.4c. However, we noticed that Map-Only takes longer to collect the first landmarks. This can be seen in Figure 8.2. We need to adjust for the outliers since participants could miss the first landmark completely even though it is very close to them. After adjusting for the outliers we can conclude that Map-Only takes a significant amount longer to collect the first landmark than VR and VR-Map, with a p-value lower than 0.005 and a positive correlation of 0.39 for the Map-Only group and negative correlations for the other two, -0.26 and -0.38 respectively. These groups collect the first landmark in around 4.4 seconds while Map-Only takes around 11.6 seconds. Why Map-Only is slower to collect the first landmark is probably because they need to adjust to the VR environment, since it is the first time they experience Holonomy VR in virtual reality. Nevertheless, we can conclude that all groups that have followed the training perform the same during the evaluation tasks which corresponds to our hypothesis.

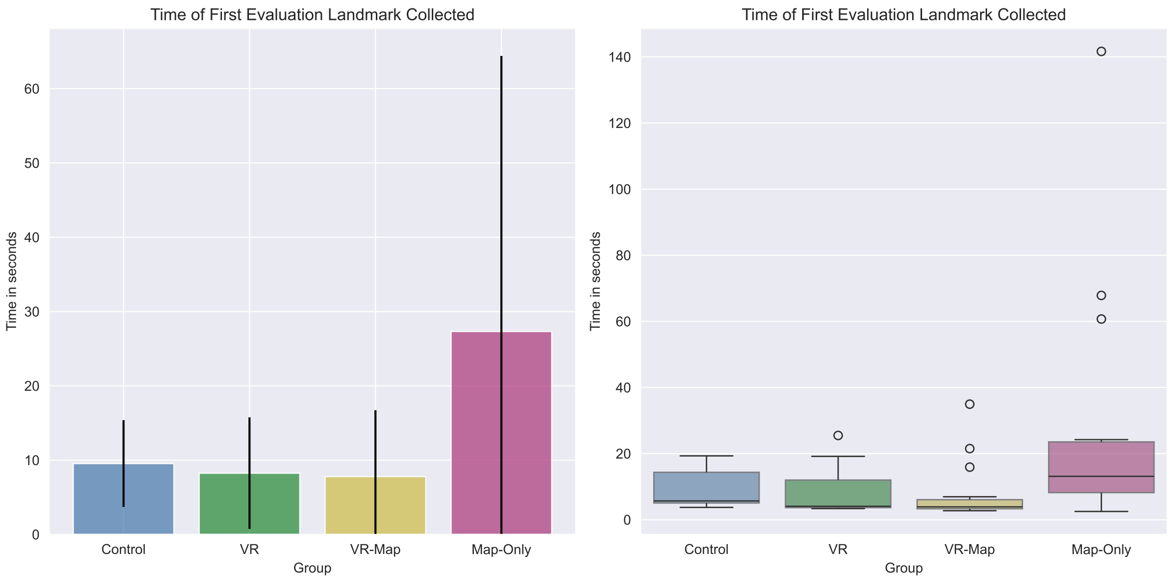

Figure 8.2: Time taken to collect the first landmark in the first evaluation

For the control group, we expect that they perform worse than the other groups since they did not have the opportunity to learn about the landmarks in the training session. For the first evaluation session, this is not the case. Here the groups perform not significantly different from each other. However, for the second evaluation session, this story is different. Here it can be seen that the control group takes significantly more steps compared to VR and Map-Only, but not to VR-Map. This might be because VR-Map was still adapting to rely only on the beacon and was focusing more on the map during practice and training. The time taken is not significantly different between groups, but these p-values are close to 0.05. This is also reflected by the moderately positive correlation of the control (0.39) and the negative correlation of VR (-0.41). The reason why the first evaluation stage does not reflect this is probably due to the fact that this stage is easier since the first two landmarks can be reached directly without the use of the beacon. The second evaluation stage is more difficult since none of the landmarks can be directly reached. However, we cannot say for certain that the lack of training is the main issue, since the control had also spent less time in Holonomy VR before the evaluation sessions.

We can now also discuss the results of the landmark questions. The results can be found in Table 7.2. This is less about the navigational performance, but it does indicate the ability to remember the location of the landmarks. We would expect the control group to perform the worst since they did not have the opportunity to learn about all the landmarks in the training. We compare the results between the groups and check if the result could not be caused by guessing. This is done by computing the binomial probability. If the p-value is less than 0.05 we say that the result is significant and that the participants have learned about this hyperbolic property.

From the results, only landmark questions 1 and 3 look promising, since the other questions were not significantly different from guessing. Landmark 1 should have been easy since the well and mushroom were placed very close together. However, only VR performed significantly better than guessing. This might be because the questions were not stated clearly enough and the participant might have thought to choose the furthest landmark from the origin. However, when it did click control performed better than guessing which increased its mean from 37.50% to 71.43% and also has a significant difference with a p-value of 0.01. This allowed the control to actually find all the landmarks, but when it did not click they could not know the location of these landmarks. This should mean however that the questions stated were clear, otherwise, the result of the control group should also indicate that it could have been guessing.

The other landmark question that is useful is landmark question three. This question was relatively simple since the bee and chest are close to each other and the golem is further away. Every group performed better than guessing. However, group VR-Map performed significantly worse than VR and Map-Only which might indicate that group VR-Map was more focused on the map than being aware of its surroundings. The control also performs worse but not significantly. However, this question could also be answered by an educated guess since, from these three landmarks, the golem is the only landmark that is not close to the origin.

To find out if the question was actually understood correctly we can look at the second question. No group answered this question significantly better than guessing, however, when we select the furthest landmark from the origin as the correct answer every group answers this question correctly. This suggests that the question was not clear and participants thought that they should select the landmark which was the furthest away from the origin.

So landmarks can be remembered but the question was not clear and thus we cannot make clear conclusions about the cognitive map of the participants. However, what we can conclude is that reaching all the landmarks allows the participant to give a better answer to the questions which makes sense since the participant otherwise would not know about the existence of the landmark.

For the fourth hypothesis we need to compare our findings to the findings about navigation on a Euclidean plane. Since the groups that followed the training perform the same in the evaluation session, we can conclude that this corresponds to results achieved in studies about navigation on a Euclidean plane. However, when looking at the training times, we only found that Map-Only performed faster than the other groups during training. This indicates that in the VR experience the access to a map does not improve the training time compared to the usage of only the beacon. So we can conclude that purely using a map is a faster training method in hyperbolic space, which corresponds to the research about navigation on a Euclidean plane. However, the current cues available in the VR experience do not impact the training time, which does not correspond to the research about navigation on a Euclidean plane.

Although not related to the hypotheses, the SBSOD score of the participants might have an impact on the understanding of the beacon cue. The SBSOD score relates to the self-rated sense of direction of the participant. From the data, which can be found in Table A.4, it can be seen that the SBSOD score is significantly higher when it clicked for the VR group with a p-value of 0.001 and a positive correlation of 0.75. This might be caused by the fact that VR needed to purely rely on the beacon. Although this is also the case for the control group, they did spend less time in total with the beacon. So by having a good sense of direction, the beacon cue was easier to interpret.

Another thing that is worth while to discuss is the different strategies that were observed[1]. The successful strategies were mostly related to the beacon. Participants who kept looking at the beacon figured out with trial and error how to get the beacon closer. Another good strategy is to get as close to the beacon as possible and walk around the closest tree until the world has rotated in such a way that you can get closer to the beacon again. Some participants figured out what the optimal direction to walk around the tree was, others tried until they could walk closer. However, some participants stopped after only one rotation which might not be enough. Then they figured that they were doing something wrong when they actually were not.

A few other unsuccessful strategies that were observed were random walking, going around all four trees, and walking around the tree furthest from the beacon. Going around all the trees does allow you to easily backtrack your steps, however, one mistake and you become easily lost. Walking around the tree furthest from the beacon rotates the world in the same way as the closest tree, but you cannot walk closer to the beacon which does not allow you to progress.

In the next section, we discuss the two hypotheses that remain. These are related to the hyperbolic properties.

[1]The strategies were not specifically noted down during the user study.

8.2. Hyperbolic Properties

We start with the fifth hypothesis. This hypothesis is about learning the affordances that the hyperbolic environment presents and to see if the participant can relate them back to the properties of the hyperbolic environment. The results for each group can be found in Table 7.3. There are a total of 4 questions that relate to the different properties. As with the discussion about the landmark questions, we will compare the results between the groups and check if the result could not be caused by guessing.

First, we start with the question about the number of tiles. This question asks how many tiles there are around a single point. This question was only answered 22.58% correctly by all participants. This is not significantly different between the groups and might be the result of guessing. However, we could score this question better by taking into account that 5 and 6 tiles around a point are hyperbolic, 4 tiles are Euclidean and 2 and 3 tiles are spherical. Scoring spherical geometry as -1, Euclidean as 0 and hyperbolic as 1 we get a mean of 27.41% which also does not display significant differences between the groups or the old method and it could still be caused by guessing. From this, we can conclude that the participants unfortunately did not learn about the property that more squares can fit around a single point in hyperbolic geometry compared to Euclidean geometry.

Next, we take a look at the question about ending up in the same location. This question asked if you go one tile forward and then one tile to the right, do you end up in the same location as if you go one tile left and one tile forward. This question was answered 88.71% correctly by all participants. This is significantly different from guessing with a p-value of 0.000. Another thing to note is that the Map-Only group answers this question always right. Although this is only significantly different from the control with a p-value of 0.039 it is still good to note. For the control group, this is probably because they spent less time in hyperbolic space since they did not complete the training, this is also reflected in the high standard deviation. The reason why Map-Only did better than the other groups might be because they could immediately see what their actions implied and they could easily see that they would not end up in the same location. However, also VR and VR-Map scored high on this question which also implies that they realised this eventually. So we can conclude that the participants have learned about the property that taking mirrored paths with a 90-degree turn does not end up in the same location on the hyperbolic plane.

The third question is about the behaviour of two paths that start in the same direction. This question asked if two paths that start in the same direction will cross, run parallel or will diverge. This question differed significantly between the groups. The control and Map-Only group scored significantly better than guessing but VR and VR-Map did not. Since the control group still has a large standard deviation and only has a score of 62% we will assume that only Map-Only has actually learned about this property. This does make sense since the Map-Only group could easily see that paths diverged in the 2D experience. However, the VR-Map had also access to the map, but could not see the bushes and trees on it which might explain this difference. If this is true it would mean that the representation of the bushes and trees on the map is important to learn about this property.

Another thing to note is that this result is moderately correlated with studying ”Information and Communication Technologies” and ”Social sciences, journalism and information” with a positive and negative correlation of 0.3 respectively. This difference is also significant with a p-value of 0.002. This might be because the first study is more technical and the second study is more about the social aspects of information. This result is also moderately correlated with the training completion ratio, however, this is most likely because all participants in the Map-Only group have completed the training. From these results, we can conclude that the Map-Only group has learned about the property that two paths that start in the same direction will diverge on the hyperbolic plane.

The last question is about the rotation of the world when visiting all tiles around a tree. None of the groups scored significantly better than guessing. However, when we only take the participants into account where it clicked, then it is a different story. Now the control group and VR-Map score significantly better than guessing with a p-value of 0.045 and 0.039 respectively they achieved a mean of 71.43% and 63.64% respectively. This might be because the control group had less time spent in Holonomy VR before the training and thus needed to realise better how the beacon rotated around them and the same holds for VR-Map since in the evaluation they needed to rely on the beacon completely for the first time. When closely paying attention to the beacon this property can also be quickly seen. So we can conclude that for the participants that clicked and spent less time using the beacon before the evaluation session they have learned about the property that the world rotates clockwise when visiting all tiles around a tree. This might indicate that participants that spent more time with the beacon were relying more on their intuition than on the behaviour of the beacon.

So yes, participants can relate the affordances of the hyperbolic environment to the properties of hyperbolic space. This also differs between the 2D and VR experience. However, we cannot say that participants learn about the properties because of the training method or the experience of all sessions combined.

For the last and sixth hypothesis we hypothesised that the understanding of hyperbolic properties improves navigational performance in hyperbolic space. This is only partly true since we only see an improvement over guessing in the control and Map-Only groups when it clicked and asked about the rotation of the world. The difference between clicked and not clicked is, however, not significant, but it does show that participants who understood how to navigate performed better in these groups. So we can only conclude that the understanding of this specific hyperbolic property improves the ability to navigate in hyperbolic space.

Finally, we found that most participants who understood how to navigate enjoyed the experience. To quantify this statement, we did a sentiment analysis using DistilBERT [47] which is based on BERT [11]. This model can process natural language and classifies it as positive or negative. When it clicked, participants had an average positive sentiment of 74.91% versus an average of 49.90% when it did not. This is a significant difference with a p-value of 0.037 and a low positive correlation of 0.27. The result is mostly affected by the control group where there is a significant difference with a p-value of 0.023 and a moderate positive correlation of 0.56. This is probably because the control group had less time to enjoy the experience, which would make it even more frustrating when it did not click.

9 Conclusion

Now that we have analysed the sub-research questions we can answer the main research question. The main research question is "How does navigational training in a hyperbolic environment influence navigational performance in this environment and the understanding of its hyperbolic properties?" We hypothesised in our main hypothesis that navigational training in a hyperbolic environment will improve navigational performance in this environment and the understanding of its hyperbolic properties.

First, we look at the first part of this statement. Does navigational training in a hyperbolic environment improve navigational performance in this environment? We cannot state this for certain. From the first evaluation, we can only state that this is not the case, since all groups performed the same. However, this is also the easiest evaluation session. From the second evaluation session, it is clear that the control group performed worse than VR and Map-Only, however not significantly worse than VR-Map. This could be because of the lack of navigational training or because the group did not have enough practice time to efficiently navigate to begin with. Unfortunately, we cannot say which is the case. We could have improved this by letting the control group practise during the training by finding other landmarks or completing another game mode offered by Holonomy VR.

What we can say is that participants can adapt to navigating in a confined hyperbolic space and that all cues used were effective to get familiar with Holonomy VR. Especially the 2D experience allowed participants to get quickly familiar with the environment and its behaviour. Most participants could apply what they have learned in the 2D experience to the actual VR experience. We think that it is possible to increase the number of completed stages by giving the participants more instructions. This would probably have helped with the adaptation process and allowed more participants to navigate effectively. However, this would have affected the learning about hyperbolic properties, since these instructions would explain more about the actual hyperbolic world.

Next, does navigational training in hyperbolic space improve the understanding of the hyperbolic properties? The questions that were asked in the survey touched four different properties. From the question about the number of tiles around a tree, we can conclude that participants did not realise how many tiles there actually are and thus not realise that more space can fit around a single point in hyperbolic space compared to Euclidean space. The next question asked about the mirrored paths with a 90-degree turn, which will end up in the same location on a Euclidean plane but not on a hyperbolic plane. Most participants answered this question correctly and thus realised that this is not the case in hyperbolic geometry. The third question asked about two paths that start in the same direction and diverge on the hyperbolic plane. Only the participants that had the 2D experience answered this question correctly. This means that the 2D experience is effective in teaching this property. The last question asked about the rotation of the world, which corresponds to the holonomy property. Only the participants that had spent less time using the beacon before the evaluation, because they did not do the training, or had access to the map on their wrist, and could effectively navigate the environment answered this question correctly. This means that the beacon could be effective in teaching this property, however, when intuition takes over participants do not actively recall this property.

So yes, people can actually learn about some of the properties of hyperbolic space by navigating in a hyperbolic environment. We can also conclude that even the control group learned about the properties and has thus spent enough time in the environment. The training type also influences the result since the 2D experience allows the participants to learn about the diverging paths property.

To conclude, we say that people can learn by themselves how to navigate in a hyperbolic environment and learn about some of the properties of hyperbolic space. However, we cannot state that navigational training by finding landmarks is effective in improving the navigational performance. However, it is clear that people can learn about the properties of hyperbolic space by navigating in a hyperbolic environment and that the training type influences this result.

In this thesis, we have taken an extensive look into navigating in a hyperbolic environment. Now that we know that people can navigate in hyperbolic space we could explore the use of the technique in other applications to improve issues with motion sickness or other types of cybersickness. Currently, cybersickness was not measured but most people were comfortable with the headset and the experience even though they spent a lot of time in virtual reality. This indicates that this method could be worthwhile to explore further.

Another thing that we have learned is that participants that grasped the concept found it fun and interesting to take part in the study and some even wanted to learn more about the subject. This could make it worth while to investigate the immersion and entertainment value of the environment in the future. To make the setup more flexible we could explore the use of different tilings of the hyperbolic plane and how this influences the results. This would require an adaptation to the generation of the map and to the rendering of the tiles in the virtual environment to allow the adaptation to different tilings. Another thing that could be explored is to enlarge the current 3x3 grid, this would give the participant more space to move around and be able to reach destinations faster.

References

[1] 3dfancy. *Lowpoly Wells*. Aug. 2016.

[2] amusedART. *Fantasy Bee*. Dec. 2019.

[3] Alaa eddine Bensad and Aziz Ikemakhen. "Hyperbolic Barycentric Coordinates and Applications". In: *Computer Aided Geometric Design* 95 (May 2022), p. 102086. ISSN: 0167-8396. DOI: 10.1016/j.cagd.2022.102086.

[4] Otmar Bock, Sylvie Abeele, and Udo Eversheim. "Human Adaptation to Rotated Vision: Interplay of a Continuous and a Discrete Process". In: *Experimental Brain Research* 152.4 (Oct. 2003), pp. 528–532. ISSN: 1432-1106. DOI: 10.1007/s00221-003-1643-x.

[5] Maksim Bugrimov. *Magic Idol*. May 2015.

[6] Stefano Burigat and Luca Chittaro. "Navigation in 3D Virtual Environments: Effects of User Experience and Location-Pointing Navigation Aids". In: *International Journal of Human-Computer Studies* 65.11 (Nov. 2007), pp. 945–958. ISSN: 1071-5819. DOI: 10.1016/j.ijhcs.2007.07.003.

[7] Polona Caserman et al. "Cybersickness in Current-Generation Virtual Reality Head-Mounted Displays: Systematic Review and Outlook". In: *Virtual Reality* 25.4 (Dec. 2021), pp. 1153–1170. ISSN: 1434-9957. DOI: 10.1007/s10055-021-00513-6.

[8] Edgar Chan et al. "From Objects to Landmarks: The Function of Visual Location Information in Spatial Navigation". In: *Frontiers in Psychology* 3 (Aug. 2012). ISSN: 1664-1078. DOI: 10.3389/fpsyg.2012.00304.

[9] Cmglee. *Comparison of Elliptic, Euclidean and Hyperbolic Geometries in Two Dimensions*. Oct. 2020.

[10] CodeParade. *Hyperbolica on Steam*. https://store.steampowered.com/app/1256230/Hyperbolica/.

[11] Jacob Devlin et al. *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. May 2019. DOI: 10.48550/arXiv.1810.04805. arXiv: 1810.04805 [cs].

[12] Dooder. *Honey Insect Wing Buzz - Weather & Seasons Icons*.

[13] Dungeon Mason. *RPG Monster Buddy*. May 2023.

[14] Russell A. Epstein et al. "The Cognitive Map in Humans: Spatial Navigation and Beyond". In: *Nature Neuroscience* 20.11 (Nov. 2017), pp. 1504–1513. ISSN: 1546-1726. DOI: 10.1038/nn.4656.

[15] Micheal Erickson. *The Change Process*. Nov. 2013.

[16] Martin J. Farrell et al. "Transfer of Route Learning From Virtual to Real Environments". In: *Journal of Experimental Psychology: Applied* 9.4 (2003), pp. 219–227. ISSN: 1939-2192. DOI: 10.1037/1076-898X.9.4.219.

[17] James Jerome Gibson. *The Ecological Approach to Visual Perception*. Psychology Press, 1986. ISBN: 978-0-89859-959-6.

[18] Lara Gregorians and Hugo J. Spiers. "Affordances for Spatial Navigation". In: *Affordances in Everyday Life: A Multidisciplinary Collection of Essays*. Ed. by Zakaria Djebbara. Cham: Springer International Publishing, 2022, pp. 99–112. ISBN: 978-3-031-08629-8. DOI: 10.1007/978-3-031-08629-8_10.

[19] Vi Hart et al. *Non-Euclidean Virtual Reality I: Explorations of H3*. Feb. 2017. arXiv: 1702.04004 [math].

[20] Vi Hart et al. *Non-Euclidean Virtual Reality II: Explorations of H2xE*. Feb. 2017. arXiv: 1702.04862 [math].

[21] Robin Hartshorne. *Geometry: Euclid and Beyond*. Ed. by S. Axler, F. W. Gehring, and K. A. Ribet. Undergraduate Texts in Mathematics. New York, NY: Springer, 2000. ISBN: 978-1-4419-3145-0 978-0-387-22676-7. DOI: 10.1007/978-0-387-22676-7.

[22] Mary Hegarty et al. "Development of a Self-Report Measure of Environmental Spatial Ability". In: *Intelligence* 30.5 (Sept. 2002), pp. 425–447. ISSN: 0160-2896. DOI: 10.1016/S0160-2896(02)00116-2.

[23] Lukas Hejtmanek et al. "How Much of What We Learn in Virtual Reality Transfers to Real-World Navigation?" In: *Multisensory Research* (Mar. 2020). DOI: 10.1163/22134808-20201445.

[24] *Hyperbolic Geometry*. Springer Undergraduate Mathematics Series. London: Springer-Verlag, 2005. ISBN: 978-1-85233-934-0. DOI: 10.1007/1-84628-220-9.

[25] Kevin Iglesias. *Giant Monster Model - Golem*. Aug. 2024.

[26] ithappy. *Low Poly 3D Models Pack*. Jan. 2025.

[27] Todd B. Kashdan et al. "The Curiosity and Exploration Inventory-II: Development, Factor Structure, and Psychometrics". In: *Journal of Research in Personality* 43.6 (Dec. 2009), pp. 987–998. ISSN: 0092-6566. DOI: 10.1016/j.jrp.2009.04.011.

[28] Numair Khan and Anis Ur Rahman. "Rethinking the Mini-Map: A Navigational Aid to Support Spatial Learning in Urban Game Environments". In: *International Journal of Human–Computer Interaction* 34.12 (Dec. 2018), pp. 1135–1147. ISSN: 1044-7318. DOI: 10.1080/10447318.2017.1418804.

[29] ROBERT M. Kitchin. "Methodological Convergence in Cognitive Mapping Research: Investigating Configurational Knowledge". In: *Journal of Environmental Psychology* 16.3 (Sept. 1996), pp. 163–185. ISSN: 0272-4944. DOI: 10.1006/jevp.1996.0015.

[30] Sebastian Koenig et al. "Assessing Navigation in Real and Virtual Environments: A Validation Study". In: *International Journal on Disability and Human Development* 10.4 (Nov. 2011), pp. 325–330. ISSN: 2191-0367. DOI: 10.1515/IJDHD.2011.050.

[31] Eryk Kopczyński, Dorota Celińska-Kopczyńska, and Marek Čtrnáct. "HyperRogue: Playing with Hyperbolic Geometry". In: July 2017.

[32] Javier Marín-Morales et al. "Navigation Comparison between a Real and a Virtual Museum: Time-dependent Differences Using a Head Mounted Display". In: *Interacting with Computers* 31 (July 2019). DOI: 10.1093/iwc/iwz018.

[33] Martin Lu. *Hyperbolic Geometry for the Uninitiated and Curious- How to Simulate Hyperbolic Space*. July 2020.

[34] Jesus Mayor et al. "Virtual Reality Presence in Partially Non-Euclidean Environments". In: *PRESENCE: Virtual and Augmented Reality* (Sept. 2024), pp. 1–14. ISSN: 1054-7460. DOI: 10.1162/pres_a_00419.

[35] Ascher K. Munion et al. "Gender Differences in Spatial Navigation: Characterizing Wayfinding Behaviors". In: *Psychonomic Bulletin & Review* 26.6 (Dec. 2019), pp. 1933–1940. ISSN: 1531-5320. DOI: 10.3758/s13423-019-01659-w.

[36] Suzanne A. E. Nooij et al. "Vection Is the Main Contributor to Motion Sickness Induced by Visual Yaw Rotation: Implications for Conflict and Eye Movement Theories". In: *PLOS ONE* 12.4 (Apr. 2017), e0175305. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0175305.

[37] Avi Parush, Shir Ahuvia, and Ido Erev. "Degradation in Spatial Knowledge Acquisition When Using Automatic Navigation Systems". In: *Spatial Information Theory*. Ed. by Stephan Winter et al. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2007, pp. 238–254. ISBN: 978-3-540-74788-8. DOI: 10.1007/978-3-540-74788-8_15.

[38] Avi Parush and Dafna Berman. "Navigation and Orientation in 3D User Interfaces: The Impact of Navigation Aids and Landmarks". In: *International Journal of Human-Computer Studies* 61.3 (Sept. 2004), pp. 375–395. ISSN: 1071-5819. DOI: 10.1016/j.ijhcs.2003.12.018.

[39] Aftab E. Patla. "Understanding the Roles of Vision in the Control of Human Locomotion". In: *Gait & Posture* 5.1 (Feb. 1997), pp. 54–69. ISSN: 0966-6362. DOI: 10.1016/S0966-6362(96)01109-5.

[40] *Paul Tol's Notes*. https://sronpersonalpages.nl/~pault/.

[41] Vincent Adelizzi Pisani et al. "Navigation by Walking In Hyperbolic Space Using Virtual Reality". In: *Extended Abstracts of the Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts*. Barcelona Spain: ACM, Oct. 2019, pp. 611–618. ISBN: 978-1-4503-6871-1. DOI: 10.1145/3341215.3356287.

[42] Roos Pot-Kolder et al. "Anxiety Partially Mediates Cybersickness Symptoms in Immersive Virtual Reality Environments". In: *Cyberpsychology, Behavior and Social Networking* 21.3 (Mar. 2018), pp. 187–193. ISSN: 2152-2723. DOI: 10.1089/cyber.2017.0082.

[43] "Hyperbolic Geometry". In: *Foundations of Hyperbolic Manifolds*. Ed. by John G. Ratcliffe. New York, NY: Springer, 2006, pp. 54–99. ISBN: 978-0-387-47322-2. DOI: 10.1007/978-0-387-47322-2_3.

[44] Sharif Razzaque, Zachariah Kohn, and Mary C. Whitton. *Redirected Walking*. 2001. DOI: 10.2312/egs.20011036.

[45] Sharif Razzaque et al. "Redirected Walking in Place". In: *Eurographics Workshop on Virtual Environments*. The Eurographics Association, 2002, 8 pages. ISBN: 978-1-58113-535-0. DOI: 10.2312/EGVE/EGVE02/123-130.

[46] Anthony E. Richardson, Daniel R. Montello, and Mary Hegarty. "Spatial Knowledge Acquisition from Maps and from Navigation in Real and Virtual Environments". In: *Memory & Cognition* 27.4 (July 1999), pp. 741–750. ISSN: 1532-5946. DOI: 10.3758/BF03211566.

[47] Victor Sanh et al. *DistilBERT, a Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter*. Mar. 2020. DOI: 10.48550/arXiv.1910.01108. arXiv: 1910.01108 [cs].

[48] L. Simón-Vicente et al. "Cybersickness. A Systematic Literature Review of Adverse Effects Related to Virtual Reality". In: *Neurología* 39.8 (Oct. 2024), pp. 701–709. ISSN: 0213-4853. DOI: 10.1016/j.nrl.2022.04.009.

[49] SineVFX. *Translucent Crystals*. May 2025.

[50] Martin Skrodzki et al. *Holonomy: A Virtual Reality Exploration of Hyperbolic Geometry*. Sept. 2024. DOI: 10.1145/3665318.3677149.

[51] Dajana Snopková et al. "Navigation in Indoor Environments: Does the Type of Visual Learning Stimulus Matter?" In: *ISPRS International Journal of Geo-Information* 8.6 (June 2019), p. 251. ISSN: 2220-9964. DOI: 10.3390/ijgi8060251.

[52] Frank Steinicke et al. "Analyses of Human Sensitivity to Redirected Walking". In: *Proceedings of the 2008 ACM Symposium on Virtual Reality Software and Technology*. VRST '08. New York, NY, USA: Association for Computing Machinery, Oct. 2008, pp. 149–156. ISBN: 978-1-59593-951-7. DOI: 10.1145/1450579.1450611.

[53] Parcly Taxel. *Order-5 Square Tiling. Vector Version of H2 Tiling 245-4*. Feb. 2020.

[54] Fred the Oyster. *Intrinsic Curvature of the Sphere*. Sept. 2014.

[55] UNESCO Institute for Statistics. *International Standard Classification of Education: Fields of Education and Training 2013 (ISCED-F 2013) Detailed Field Descriptions*. UNESCO Institute for Statistics, 2015. ISBN: 978-92-9189-179-5. DOI: 10.15220/978-92-9189-179-5-en.

[56] Gerard Venema. "The Poincaré Disk". In: *Exploring Advanced Euclidean Geometry with GeoGebra*. Vol. 44. Classroom Resource Materials. American Mathematical Society, 2013, pp. 96–106. ISBN: 978-0-88385-784-7 978-1-61444-111-3 978-1-4704-6229-1. DOI: 10.5948/9781614441113.

[57] William H. Warren et al. "Wormholes in Virtual Space: From Cognitive Maps to Cognitive Graphs". In: *Cognition* 166 (Sept. 2017), pp. 152–163. ISSN: 0010-0277. DOI: 10.1016/j.cognition.2017.05.020.

[58] Séamas Weech, Sophie Kenny, and Michael Barnett-Cowan. "Presence and Cybersickness in Virtual Reality Are Negatively Related: A Review". In: *Frontiers in Psychology* 10 (Feb. 2019). ISSN: 1664-1078. DOI: 10.3389/fpsyg.2019.00158.

[59] Robert B. Welch. *Perceptual Modification: Adapting to Altered Sensory Environments*. Elsevier, 1978. ISBN: 978-1-4832-7478-2.

[60] Jan M. Wiener, Simon J. Büchner, and Christoph Hölscher. "Taxonomy of Human Wayfinding Tasks: A Knowledge-Based Approach". In: *Spatial Cognition & Computation* 9.2 (May 2009), pp. 152–165. ISSN: 1387-5868. DOI: 10.1080/13875860902906496.

[61] Jan M. Wiener et al. "A Novel Virtual-Reality-Based Route-Learning Test Suite: Assessing the Effects of Cognitive Aging on Navigation". In: *Behavior Research Methods* 52.2 (Apr. 2020), pp. 630–640. ISSN: 1554-3528. DOI: 10.3758/s13428-019-01264-8.

[62] Baran Yarar et al. *"Holonomy": A Non-Euclidean Labyrinth Game in Virtual Reality*. 2022.

A
Extra Results

Group	P Mean	P SD	T Mean	T SD	E Mean	E SD	Total Mean	Total SD
Control	72.00%	31.00%	-	-	68.00%	28.00%	69.00%	27.00%
VR	77.00%	32.00%	76.00%	27.00%	74.00%	32.00%	75.00%	25.00%
VR-Map	72.00%	36.00%	73.00%	29.00%	78.00%	28.00%	75.00%	26.00%
Map-Only	83.00%	24.00%	100.00%	0.00%	77.00%	24.00%	88.00%	10.00%

Table A.1: Percentage of completion by session

Group	P 2 Mean	P 2 SD	P 3 Mean	P 3 SD	T Mean	T SD	E 1 Mean	E 1 SD	E 2 Mean	E 2 SD
Control	89.14	71.08	71.06	56.14	-	-	162.46	68.30	255.46	62.69
VR	118.43	81.80	81.82	78.16	446.93	147.91	151.61	102.21	150.05	68.00
VR-Map	90.69	60.45	94.51	77.56	431.01	87.62	143.26	75.76	211.43	68.67
Map-Only	71.05	53.90	74.48	81.27	154.97	147.81	212.25	82.30	193.72	50.13

Table A.2: Time taken by stage

Group	P 2 Mean	P 2 SD	P 3 Mean	P 3 SD	T Mean	T SD	E 1 Mean	E 1 SD	E 2 Mean	E 2 SD
Control	27.60	24.50	36.23	27.90	-	-	90.50	36.58	169.33	29.51
VR	45.08	42.19	35.64	33.27	265.67	53.00	84.50	51.59	103.71	42.08
VR-Map	30.55	28.74	40.85	36.02	222.75	64.01	91.18	57.22	142.00	51.09
Map-Only	69.23	54.07	71.62	67.57	197.07	287.32	91.62	36.39	125.00	21.41

Table A.3: Steps taken by stage

[1]One participant left out one of the questions in the SBSOD survey, we filled this missing value with the mean of the possible options.

Group	SBSOD Clicked Mean	SBSOD Clicked SD	SBSOD Non-Clicked Mean	SBSOD Non-Clicked SD
Control	66.53%	21.28%	59.79%	18.49%
VR	68.78%	13.53%	44.29%	6.64%
VR-Map	58.27%	14.54%	61.52%	18.34%
Map-Only	52.12%	21.33%	60.23%	25.40%

Table A.4: SBSOD score in percentage when it clicked and when it did not click. A higher score indicates a self reported better sense of direction.[1]

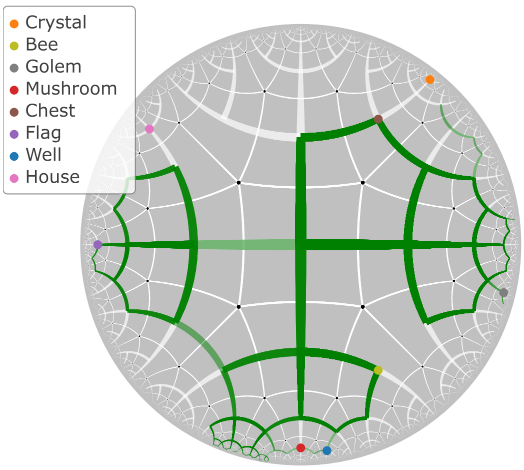

Figure A.1: Path taken by participant 36 during training. The path shows that they did not reach all the landmarks, but this is not necessary for group D to complete the training. We can see in the bottom left that they got lost and took many steps to find their way back.

B

Survey

ID: Enter your participant ID

____

Age: What is your age?

____

Gender: What is your gender?

- ☐ Male
- ☐ Female
- ☐ Non-binary / third gender
- ☐ Prefer not to say

Study: What is your field of study?

- ☐ Education
- ☐ Arts and humanities
- ☐ Social sciences, journalism and information
- ☐ Business, administration and law
- ☐ Natural sciences, mathematics and statistics
- ☐ Information and Communication Technologies
- ☐ Engineering, manufacturing and construction
- ☐ Agriculture, forestry, fisheries and veterinary
- ☐ Health and welfare
- ☐ Services
- ☐ Literacy
- ☐ Other

Attainment: What is your educational attainment?

- ☐ High school
- ☐ Bachelor degree
- ☐ Master degree
- ☐ Doctoral degree

Hand: What is your dominant hand?
- ☐ Left
- ☐ Right

Info: For the statements below, indicate if you agree or disagree with the statement.

Games I'm experienced with video games.
- ☐ Strongly disagree
- ☐ Disagree
- ☐ Somewhat disagree
- ☐ Somewhat agree
- ☐ Agree
- ☐ Strongly agree

VR I have experience with virtual reality.
- ☐ Strongly disagree
- ☐ Disagree
- ☐ Somewhat disagree
- ☐ Somewhat agree
- ☐ Agree
- ☐ Strongly agree

Hyperbolic Geometry: I have a good understanding of hyperbolic geometry.
- ☐ Strongly disagree
- ☐ Disagree
- ☐ Somewhat disagree
- ☐ Somewhat agree
- ☐ Agree
- ☐ Strongly agree

SBSOD: The following statements ask you about your spatial and navigational abilities, preferences, and experiences. After each statement, you should circle a number to indicate your level of agreement with the statement. Circle "1" if you strongly agree that the statement applies to you, "7" if you strongly disagree, or some number in between if your agreement is intermediate. Circle "4" if you neither agree nor disagree.

	strongly agree (1)	2	3	neither agree nor disagree (4)	5	6	strongly disagree (7)
1. I am very good at giving directions.	☐	☐	☐	☐	☐	☐	☐
2. I have a poor memory for where I left things.	☐	☐	☐	☐	☐	☐	☐
3. I am very good at judging distances.	☐	☐	☐	☐	☐	☐	☐
4. My "sense of direction" is very good.	☐	☐	☐	☐	☐	☐	☐
5. I tend to think of my environment in terms of cardinal directions (N, S, E, W).	☐	☐	☐	☐	☐	☐	☐
6. I very easily get lost in a new city.	☐	☐	☐	☐	☐	☐	☐
7. I enjoy reading maps.	☐	☐	☐	☐	☐	☐	☐
8. I have trouble understanding directions.	☐	☐	☐	☐	☐	☐	☐

9. I am very good at reading maps.	☐	☐	☐	☐	☐	☐	☐
10. I don't remember routes very well while riding as a passenger in a car	☐	☐	☐	☐	☐	☐	☐
11. I don't enjoy giving directions.	☐	☐	☐	☐	☐	☐	☐
12. It's not important to me to know where I am.	☐	☐	☐	☐	☐	☐	☐
13. I usually let someone else do the navigational planning for long trips.	☐	☐	☐	☐	☐	☐	☐
14. I can usually remember a new route after I have traveled it only once.	☐	☐	☐	☐	☐	☐	☐
15. I don't have a very good "mental map" of my environment.	☐	☐	☐	☐	☐	☐	☐
8. I have trouble understanding directions.	☐	☐	☐	☐	☐	☐	☐
9. I am very good at reading maps.	☐	☐	☐	☐	☐	☐	☐
10. I don't remember routes very well while riding as a passenger in a car.	☐	☐	☐	☐	☐	☐	☐
11. I don't enjoy giving directions.	☐	☐	☐	☐	☐	☐	☐

12. It's not important to me to know where I am.	□	□	□	□	□	□	□
13. I usually let someone else do the navigational planning for long trips.	□	□	□	□	□	□	□
14. I can usually remember a new route after I have traveled it only once.	□	□	□	□	□	□	□
15. I don't have a very good "mental map" of my environment.	□	□	□	□	□	□	□

Info: Rate the statements below for how accurately they reflect the way you generally feel and behave. Do not rate what you think you should do, or wish you do, or things you no longer do. Please be as honest as possible.

Ex 1: I actively seek as much information as I can in new situations.

- □ Very slightly or not at all
- □ A little
- □ Moderately
- □ Quite a bit
- □ Extremely

Ex 2: I am the type of person who really enjoys the uncertainty of everyday life.

- □ Very slightly or not at all
- □ A little
- □ Moderately
- □ Quite a bit
- □ Extremely

Ex 3: I am at my best when doing something that is complex or challenging.

- ☐ Very slightly or not at all
- ☐ A little
- ☐ Moderately
- ☐ Quite a bit
- ☐ Extremely

Ex 4: Everywhere I go, I am out looking for new things or experiences.

- ☐ Very slightly or not at all
- ☐ A little
- ☐ Moderately
- ☐ Quite a bit
- ☐ Extremely

Ex 5: I view challenging situations as an opportunity to grow and learn.

- ☐ Very slightly or not at all
- ☐ A little
- ☐ Moderately
- ☐ Quite a bit
- ☐ Extremely

Ex 6: I like to do things that are a little frightening.

- ☐ Very slightly or not at all
- ☐ A little
- ☐ Moderately
- ☐ Quite a bit
- ☐ Extremely

Ex 7: I am always looking for experiences that challenge how I think about myself and the world.

- ☐ Very slightly or not at all
- ☐ A little
- ☐ Moderately
- ☐ Quite a bit
- ☐ Extremely

Ex 8: I prefer jobs that are excitingly unpredictable.

- ☐ Very slightly or not at all
- ☐ A little
- ☐ Moderately
- ☐ Quite a bit
- ☐ Extremely

Ex 9: I frequently seek out opportunities to challenge myself and grow as a person.

- ☐ Very slightly or not at all
- ☐ A little
- ☐ Moderately
- ☐ Quite a bit
- ☐ Extremely

Ex 10: I am the kind of person who embraces unfamiliar people, events, and places.

- ☐ Very slightly or not at all
- ☐ A little
- ☐ Moderately
- ☐ Quite a bit
- ☐ Extremely

Info: Time to do the study!

Info: For the next questions, please select the landmark that is the furthest away from the given options.

Landmarks 1: From these options, which landmark is the furthest away from the others?

☐ Well: 

☐ Mushroom: 

☒ Bee: 

Landmarks 2: From these options, which landmark is the furthest away from the others?

☐ Chest:

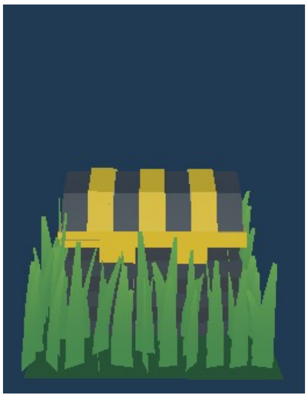

☒ Bee:

☐ Crystal:

Landmarks 3: From these options, which landmark is the furthest away from the others?

☐ Bee:

☒ Golem:

☐ Chest:

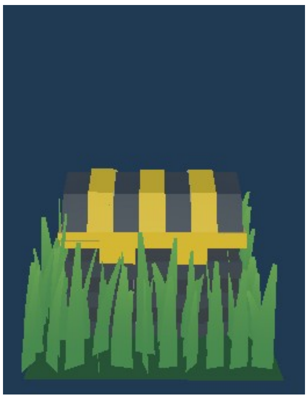

Landmarks 4: From these options, which landmark is the furthest away from the others?

☐ Flag:

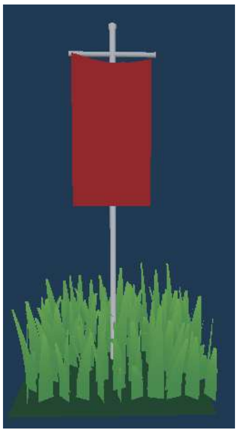

☒ Mushroom:

☐ House:

Tiles: How many tiles (squares) are there around a single tree (or point on the minimap)?

- ☐ 2 tiles
- ☐ 3 tiles
- ☐ 4 tiles
- ☒ 5 tiles
- ☐ 6 tiles

Location: When you go one tile forward, then go one tile to the left. Will you end up at the same location then if you go one tile left and one tile forward?

- ☐ Yes, I will end up at the same location
- ☒ No, I will not end up at the same location

Paths: When you take two points side by side. And start in the same direction on both of them, what happens to the two paths that you create when you go forward?

- ☐ The paths will cross:

- ☐ The paths will stay at the same distance from each other:

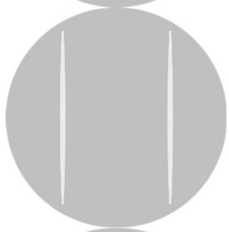

- ☒ The paths will diverge:

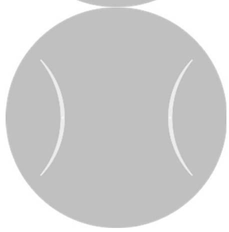

Rotation: What happens with the rotation of the world when you walk in a circle clockwise, i.e. visiting all the tiles around a tree (or point on the minimap) in the clockwise direction?

- ☒ The world will rotate 90 degrees clockwise
- ☐ The world will not rotate
- ☐ The world will rotate 90 degrees anti-clockwise

Experience: What did you think of the experience?

____

Performance: How do you evaluate your performance for the evaluation task?

____

Hyperbolic: Did you learn anything about hyperbolic geometry?

____

Extra: Do you have something else to add?

____

C

Script

Script

Group A: Control
Group B: VR
Group C: VR + Map
Group D: Map Only (Practice and training on PC, evaluation in VR)

Pre-Setup

- Do the setup on the headset (and the computer for group D)
    - Make sure that the sound is working correctly
    - Make sure that streaming is working
    - Make sure that the boundary is okay (Keep an eye on this during the user study).
    - Select the participant ID
    - Select the group
    - Select practice 1
- Open the survey in a private tab and fill in the participant ID

Give Informed Consent

Start Survey

User Study

Hints

The following hints are allowed:

- Giving hints for controls (Only applicable for group D)
- When they hear the sounds and the participant doesn't seem to get it, explain it again (Only allowed once)
- When playing in VR and a participant really seems to only wander further from the goal, tell them that they are moving away from the target.

Setup

Make sure that the Holonomy application is open and that the participant is standing at the middle of the room facing a predefined wall and reset the boundary centre by holding the "Quest Button" (most right button on the right controller). This can also be used during the study to fix boundaries for the participant.

Practice 1

[Let the participant start]

Group A and B:

Welcome to Holonomy, first we will do three practice sessions. In these sessions the goal is to find the flag. You are now in a world with trees and bushes. You can move through the world by walking. However, you will notice that you are surrounded by bushes, and it may seem that you are trapped, but this is not the case since this world is a hyperbolic world. You will notice that the world will rotate when you walk around a tree. This will move the bushes and allows you to explore different parts of the world. So please keep this in mind. When the flag is outside of the bushes, it's location will be displayed as a red beacon. When the flag enters the area inside of the bushes a positive sound will play, if it leaves the area a negative sound will play.

[Let them hear the sounds (Appear, Disappear and Collection)]

Group C:

Same as **group A and B** +

You can see a map of the world by lifting your left wrist and rotating it towards you, just like looking at a watch. You can use this map to see more of the world.

Group D:

Welcome to Holonomy, first we will do three practice sessions. In these sessions the goal is to find the flag. You now have an overview of the world. You will see some grey points and a green boundary. You can move through the world by using the **W, A, S and D** keys. You cannot move through the green boundary so it may seem that you are trapped, but this is not the case since this world is a hyperbolic world. You will notice that the world will rotate when you walk around a grey point. This will move the boundary and allows you to explore different parts of the world. So please keep this in mind. The location of the flag is indicated with a red dot.

[You can help them out with the controls]

Please walk to the flag.

[When a stage is completed, the participant can press both inner triggers two times (PC participants can press the spacebar) to go back to the menu]

Practice 2 and 3

Continue with practice 2 and then with 3.

Training

Group A:

Play pong at ponggame.org with the following settings:

Score: 20
Slow Serve: Yes
Keyboard Sensitivity: Normal
Start After Point: Alternate

Start a one player medium difficulty game with the mouse.

[Let the participant win 2 times or let the participant play 12 minutes, whichever is first completed]

[Give the explanation before starting]

Group B:

We will now start with the training. The goal is to find all the landmarks and remembering their location for the evaluation sessions. The beacons of the landmarks that you have found will turn green. The sounds that you hear only correspond to the landmarks that you have not found yet. In total you need to find eight landmarks, the order in which you find them does not matter.

Group C:

We will now start with the training. The goal is to find all the landmarks and remembering their location for the evaluation sessions. The beacons of the landmarks that you have found will turn green. The sounds that you hear only correspond to the landmarks that you have not found yet. To help you find the landmarks, you can use the map on your wrist. Here you can also see the landmarks that you need to collect. In total you need to find eight landmarks, the order in which you find them does not matter.

Group D:

We will now start with the training. The goal is to find all the landmarks and remembering their location for the evaluation sessions. To make it easier you don't have to reach the landmark but get close enough to it. This allows you to find landmarks outside of the boundary. When you are close enough the indicator on the landmark will turn green. In total you need to find eight landmarks, the order in which you find them does not matter.

Evaluation 1

[Give the explanation before starting]

Group A:

We will now start with the evaluation. The goal is to find four landmarks in a specific order.

Group B:

We will now start with the evaluation. The goal is to find four landmarks in a specific order. They are at the same location as in the training.

Group C:

Same as **Group B** +
However, you do not have access to your wrist map.

Group D:

[Move from the PC to VR]
We will now start with the evaluation. The goal is to find four landmarks in a specific order. They are at the same location as in the training. However, the points that you saw in the map have turned into trees and the boundaries you saw are now bushes. The indicators have turned into beacons. When a landmark enters the area inside of the bushes a positive sound will play, if it leaves the area a negative sound will play. The sounds that you hear only correspond to the landmarks that you have not found yet. The same principles still apply. You can move through the world by walking.
[Let them hear the sounds (Appear, Disappear and Collection)]

Evaluation 2

Continue with evaluation 2

Continue Survey

D

Informed Consent Form

Informed Consent Form

You are being invited to participate in a research study titled “Exploring Navigational Ability in Hyperbolic Space” This study is being done by Rico van Buren and Dr. Martin Skrodzki from the TU Delft and Burchin Mustafa and Prof. dr. Ineke van der Ham from Leiden University.

The purpose of this research study is to study the navigational ability when exploring a virtual environment. You will be exploring a virtual environment while being seated and while wearing a virtual reality headset with which you will be walking around. It will take you approximately 60 minutes to complete. The data will be used in the master thesis of Rico van Buren and the master thesis of Burchin Musfafa. The theses will be published by the TU Delft and Leiden University respectively. The gathered data will anonymously be published on 4TU.ResearchData to be available for further research. We will be asking you to explore virtual environments by finding specific landmarks. While exploring your performance will be tracked with the path you took and the time it took you to complete the task. After the exploration you will be asked several questions about your thoughts of the virtual environment and about the virtual environment itself. While wearing a virtual reality headset there is a possible risk of motion sickness, anxiety, claustrophobia and/or disorientation. If at any time you feel uncomfortable you can take a break or withdraw from the study completely. You will be walking within a virtual environment, to make sure you stay inside of the walkable area the boundaries of this area will be displayed by virtual bushes. The headset will also visually warn you when you get to close to the boundaries. If you accidentally leave the area, you will be guided back by one of the researchers.

As with any activity the risk of a breach is always possible. To the best of our ability your answers in this study will remain confidential. We will minimize any risks by using TU Delft and Leiden University approved software. The surveys are provided by Qualtrics and your responses to those surveys will be stored at Qualtrics. When the data is retrieved from Qualtrics it will be stored on Project Storage at TU Delft. We will be asking you about your following personal data: Age, Gender, Field of Study, Education Attainment, Dominant Hand and your experience with video games. We will not store a link between this informed consent form and the gathered data. This ensures that we gather the data anonymously. With the research data that is gathered and published you should not be identifiable.

Your participation in this study is entirely voluntary and you can withdraw at any time. You are free to omit any questions. You will not be able to withdraw your data after the session, since there is no link between you and the gathered data.

Rico van Buren
Burchin Mustafa
Dr. Martin Skrodzki
Prof. dr. Ineke van der Ham

PLEASE TICK THE APPROPRIATE BOXES	Yes	No
A: GENERAL AGREEMENT – RESEARCH GOALS, PARTICPANT TASKS AND VOLUNTARY PARTICIPATION		
1. I have read and understood the study information dated [*DD/MM/YYYY*], or it has been read to me. I have been able to ask questions about the study and my questions have been answered to my satisfaction.	☐	☐
2. I consent voluntarily to be a participant in this study and understand that I can refuse to answer questions and I can withdraw from the study at any time, without having to give a reason.	☐	☐
3. I understand that taking part in the study involves: • Exploring virtual environments while being seated. • Exploring virtual environments while wearing a virtual reality headset and walking. • Answering a survey questionnaire completed by me.	☐	☐
4. If you are a student from Leiden University and are eligible to receive study credits (EC) for participating in user studies: I understand that participating in this user study will count towards receiving study credits for participating in user studies, even if I don't complete the study. Otherwise, you can skip this question.	☐	☐
5. I understand that the study will end when the thesis projects of Rico van Buren and Burchin Mustafa are ended	☐	☐
B: POTENTIAL RISKS OF PARTICIPATING (INCLUDING DATA PROTECTION)		
6. I understand that taking part in the study involves the following risks: While wearing a virtual reality headset you can experience motion sickness, anxiety, claustrophobia and/or disorientation. I understand that these will be mitigated by the fact that I'm always able to take a break or withdraw from the session. While wearing the headset it is also possible that I accidently leave the walkable space, I understand that this space is indicated by virtual bushes and by the virtual reality headset if I come too close to the boundary. If I do leave the area I will be guided back by a researcher.	☐	☐
7. I understand that taking part in the study also involves collecting specific associated personally identifiable research data (PIRD), such as my age, my gender, my field of study, my education attainment, dominant hand and experience with video game, with the potential risk of my identity being revealed.	☐	☐
8. I understand that the following steps will be taken to minimise the threat of a data breach and protect my identity in the event of such a breach. • Use of only TU Delft and Leiden University approved software. • Surveys are provided by Qualtrics. • Responses to the surveys will be stored at Qualtrics. • Data exported from Qualtrics it will be stored on Project Storage at TU Delft. • We will create no link between directly identifiable information and the gathered research data. The following data will be published categorically: • Age (range of 5 years) • Gender • Field of Study • Education attainment • Dominant hand • Experience with video games	☐	☐

PLEASE TICK THE APPROPRIATE BOXES	Yes	No
9. I understand that personal information collected about me that can identify me, such as my name and email address will not be shared beyond the study team.	☐	☐
10. I understand that the identifiable personal data I provide will be destroyed when the study has ended.	☐	☐
C: RESEARCH PUBLICATION, DISSEMINATION AND APPLICATION		
11. I understand that after the research study the de-identified information I provide will be used for the master theses of Rico van Buren and Burchin Mustafa.	☐	☐
12. I agree that my responses, views or other input can be quoted anonymously in research outputs	☐	☐
13. I agree that any drawings that I produce as a response can be used anonymously in research outputs	☐	☐
D: (LONGTERM) DATA STORAGE, ACCESS AND REUSE		
14. I give permission for the de-identified data, such as my age, my gender, my field of study, my education attainment, dominant hand and experience with video games, that I provide to be archived in the 4TU.ResearchData repository so it can be used for future research and learning.	☐	☐

Signatures

________________________ ________________________ ________

Name of participant [printed] Signature Date

I, as researcher, have accurately read out the information sheet to the potential participant and, to the best of my ability, ensured that the participant understands to what they are freely consenting.

________________________ ____________________ ________

Researcher name [printed] Signature Date

Study contact details for further information:

Rico van Buren r.p.g.vanburen@student.tudelft.nl
Burchin Mustafa b.mustafa@umail.leidenuniv.nl
Dr. Martin Skrodzki m.skrodzki@tudelft.nl
Prof. dr. Ineke van der Ham c.j.m.van.der.ham@fsw.leidenuniv.nl